

Meshless data-driven decompositions with RBF-based inner products

M. Ratz^{a,b,*}, A. Parente^{b,d,e}, M. A. Mendez^{a,b,c}

^avon Karman Institute for Fluid Dynamics, Waterloosesteenweg 72, Sint-Genesius-Rode, Belgium

^bAero-Thermo-Mechanics Laboratory, École Polytechnique de Bruxelles, Université Libre de Bruxelles, Av. Franklin Roosevelt 50, Brussels, 1050, Belgium

^cExperimental Aerodynamics and Propulsion Lab, Universidad Carlos III de Madrid, Av. de la Universidad 30, 28911 Leganés, Spain

^dBrussels Institute for Thermal-Fluid Systems and Clean Energy (BRITE), Brussels, 1050, Belgium

^eWEL Research Institute, 1300 Wavre, Belgium

Abstract

Data-driven modal decompositions are useful tools for compressing data or identifying dominant structures. Popular ones like the dynamic mode decomposition (DMD) and the proper orthogonal decomposition (POD) are defined with continuous inner products. These are usually approximated with samples of data uniform in space and time. However, not every dataset fulfills this requirement. Numerical simulations with smoothed particle hydrodynamics or experiments with Lagrangian particle tracking velocimetry produce scattered data varying in time and space, rendering sample-based inner products impossible. In this work, we extend a previous approach that computes the modal decompositions with meshfree radial basis functions (RBFs). We regress the data and use the continuous representation of the RBFs to compute the required inner products. We choose our basis to be constant in time, greatly reducing the computational cost since the inner product of the data reduces to the inner product of the basis functions. We use this approach in the most popular decompositions, namely the POD, DMD, multi-scale POD, and the two versions of the spectral POD. For all decompositions, the RBFs give a mesh-free representation of the spatial structures. Two test cases are considered: particle image velocimetry measurements of an impinging jet and large eddy simulations of the flow past a transitional airfoil. In both cases, the RBF-based approach outperforms classical binning and better recovers relevant structures across all data densities.

Keywords: Meshless, radial basis function, modal analysis, proper orthogonal decomposition, dynamic mode decomposition, data-driven decompositions

1. Introduction

Data-driven modal decompositions such as proper orthogonal decomposition (POD) and dynamic mode decomposition (DMD) are widely used to extract dominant structures and dynamics from complex systems. These techniques have been extensively applied to obtain compact representations of high-dimensional data across diverse fields, including fluid mechanics [1–3], image processing and filtering [4], face recognition [5], reacting flows [6], structural fault detection [7] and vibration [8], biomechanics [9], neuroscience [10], epidemiology [11] to mention but a few.

Despite their wide applicability, the algorithms used to compute data-driven modal decompositions are traditionally formulated under the assumption that data are available on a spatial grid—as in image processing, experimental techniques such as particle image velocimetry (PIV, [12]), or standard numerical frameworks like finite difference, finite volume, and finite element methods [13]. This grid may be structured or unstructured, dense or sparse, but is typically assumed to be fixed in time. Such an assumption simplifies computing the inner products inherent to any linear decomposition [14], from

classical methods like POD [15–17] and DMD [2, 18] to more recent variants such as spectral POD (SPOD) [19, 20] and multi-scale POD (mPOD) [21, 22]. The spatial structures on this discrete grid also allow to link dominant features to specific physical locations.

A variant within this context is the so-called gappy scenario, where an underlying grid or spatial structure is still assumed, but the available data have values missing in space at scattered locations in time. A common strategy in this setting is to resort to interpolation [23], radial basis functions [24], or low-rank matrix completion [25] to reconstruct the missing data either prior to or during the computation of a modal decomposition. Examples of such approaches include gappy POD [26–28], the recently proposed gappy Spectral POD [29], and interpolative variants of DMD designed to handle non-uniform data in space [30] or time [31, 32]. Alternatively, when interpolation or regression methods become ineffective due to excessive gappiness, compressed sensing techniques [33, 34] can be employed to exploit sparsity and directly reconstruct the dominant flow structures from limited measurements, enabling full-field recovery from observations restricted to a sparse subset of grid points. These ideas have opened the path toward compressed sensing formulations of POD [35, 36] and DMD [34, 37].

A more significant departure from the gappy framework

*Corresponding author

Email address: manuel.ratz@vki.ac.be (M. Ratz)

arises when the data are not simply incomplete, but intrinsically *scattered*, with no underlying spatial grid or fixed set of sampling locations. This situation arises in Lagrangian computational methods, such as smoothed particle hydrodynamics [38] or vortex methods [39], where the computational particles themselves move with the flow. Similarly, in experimental techniques like Lagrangian particle tracking velocimetry [40–43], measurements are obtained at the positions of tracer particles advected by the fluid, leading to data sampled at locations varying in both space and time.

To the authors’ knowledge, there is no established framework for applying modal decompositions directly to data in their scattered, Lagrangian form. The most common strategy is to circumvent the irregularity by mapping the data onto a fixed spatial grid—typically through interpolating, projecting, or binning techniques—and then proceeding with traditional grid-based decomposition methods [44–46]. While this preprocessing step enables to use established tools like POD (e.g. [47]) and DMD (e.g. [48]), it introduces interpolation errors and may obscure important features of the original scattered data, particularly in regions with sparse sampling or rapid deformation. More advanced binning strategies iterate between mapping the data to the grid and computing the POD and progressively improve the spatial modes [49, 50].

Within the image velocimetry community, significant effort has recently been devoted to methods for mapping scattered data onto fixed grids [51–53]. More recently, however, a meshless paradigm based on radial basis function (RBF) has begun to emerge [54–59], aiming to remove grids entirely. These approaches are considered meshless since they operate directly on scattered data without requiring a structured or unstructured mesh, and do not assume a fixed spatial discretization for computing derivatives, integrals, or more complex tasks such as solving partial differential equations [24, 60–62].

Recently, the meshless RBF framework was extended to compute a POD directly from scattered data [63]. The approach is conceptually related to functional principal component analysis (FPCA) [64–66], which generalizes traditional principal component analysis (PCA) to the case where each observation is treated as a continuous function rather than a finite-dimensional vector. The key difference, among others discussed in Tirelli et al. [63], is that the FPCA uses interpolation methods to approximate an infinite-dimensional eigenvalue problem, whereas the meshless POD retains the classical POD structure but replaces the grid-based inner product with RBF-based quadrature.

This work extends the framework introduced in Tirelli et al. [63] in several important respects. First, we replace the interpolative RBF formulation with a regression-based approach and substitute the Gauss–Legendre quadrature used to assemble the temporal correlation matrix with a fully RBF-based quadrature scheme. This modification not only reduces computational cost by enabling the reuse of precomputed weight matrices but also improves flexibility in handling complex geometries. Second, we introduce an RBF-based projection strategy operating directly on the continuous

functional form obtained from the RBF approximation of the snapshots, enabling efficient evaluation of the spatial structure of each mode. Finally, the proposed framework is generalized beyond classical POD to encompass more advanced modal decompositions, including the two SPODs, mPOD, and DMD.

The remainder of the paper is organized as follows. Section 2 recalls the fundamentals of RBF regression, how they are used to compute spatial inner products, and how each of the decompositions is expressed in the RBF framework. Section 3 presents the two selected test cases and Section 4 presents the results of the meshless algorithms compared to classic binning. Finally, Section 5 closes with conclusions and perspectives.

2. Mathematical framework

The goal of all data-driven decompositions is to express data with respect to a set of basis functions. Here, we treat a two-dimensional velocity field in a domain $\Omega \in \mathbb{R}^2$, which we denote as $\mathbf{u}(\mathbf{x}, t) = (u(\mathbf{x}, t), v(\mathbf{x}, t))$, with $\mathbf{x} = (x, y) \in \Omega$ an (arbitrary) point in space and $t \in [0, T]$ an (arbitrary) point in time. The approach readily extends to higher dimensions and to scalar quantities such as temperature or density.

The general modal decomposition, in a continuous sense, reads

$$\mathbf{u}(\mathbf{x}, t) = \sum_r^{n_R} \sigma_r \boldsymbol{\phi}_r(\mathbf{x}) \psi_r(t), \quad (1)$$

where n_R is the number of modes, $\boldsymbol{\phi}_r(\mathbf{x})$ is a vector valued basis function for the spatial domain, and $\psi_r(t)$ is a scalar basis function for the time domain. These bases have unitary norm with respect to a certain measure in the spatial and temporal domain, such that the scalars σ_r are amplitudes of the modal contribution. The inner product defining these norms and the associated projection operator are

$$\langle p(t), q(t) \rangle_T = \frac{1}{T} \int_0^T q^*(t) p(t) dt \quad (2)$$

for two continuous scalar functions $p(t), q(t)$ in $t \in [0, T]$ and

$$\langle \mathbf{p}(\mathbf{x}), \mathbf{q}(\mathbf{x}) \rangle_\Omega = \frac{1}{|\Omega|} \int_\Omega \mathbf{q}^*(\mathbf{x}) \mathbf{p}(\mathbf{x}) d\mathbf{x}, \quad (3)$$

for two continuous vector fields $\mathbf{p}(\mathbf{x}), \mathbf{q}(\mathbf{x})$ in $\mathbf{x} \in \Omega$, with $|\Omega|$ a measure of the domain (an area in 2D or a volume in 3D). Throughout the paper, we use the asterisk superscript to denote both a complex conjugate of a scalar or the Hermitian transpose of a vector or matrix.

All traditional grid-based decompositions use approximations of the integrals in (2) and (3) [14]. In the simplest case of uniform sampling in time, such that $\mathbf{t} = [\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_{n_t}] = \{\mathbf{t}_i\}_{i=1}^{n_t}$ is a vector collecting the equally distributed times $\mathbf{t}_i$ and $T = n_t \Delta t$, the samples of the functions $p(\mathbf{t}_i), q(\mathbf{t}_i)$ can be stored in vectors $\mathbf{p}, \mathbf{q} \in \mathbb{R}^{n_t}$ and the inner product product (2) can be approximated as:

$$\langle p(t), q(t) \rangle_T \approx \langle p(\mathbf{t}_i), q(\mathbf{t}_i) \rangle_{T,d} = \frac{1}{n_t} \mathbf{q}^* \mathbf{p}, \quad (4)$$

with $\langle \cdot \rangle_{T,d}$ denoting the discrete approximation of the continuous inner product.

Similarly, in the spatial domain, denoting as $\{\mathbf{X}_n\}_{n=1}^{n_s}$ the set of n_s sample locations in space, and assuming that each of these is associated to the same partitions of the domain $\Delta\Omega_i$ (an area in 2D or a volume in 3D), and denoting as $\mathbf{p}, \mathbf{q} \in \mathbb{R}^{2n_s}$ the flattened matrix collecting the samples $\mathbf{p}(\mathbf{X}_i), \mathbf{q}(\mathbf{X}_i)$, (3) can be approximated as

$$\langle \mathbf{p}(\mathbf{x}), \mathbf{q}(\mathbf{x}) \rangle_{\Omega} \approx \langle \mathbf{p}(\mathbf{X}_i), \mathbf{q}(\mathbf{X}_i) \rangle_{\Omega,d} = \frac{1}{n_s} \mathbf{q}^* \mathbf{p}, \quad (5)$$

having introduced the discrete spatial inner product. These discrete inner products (4) and (5) can be easily adjusted with weight matrices for non-uniform sampling. In the case of scattered data in space or time, the approximations (4) and (5) do not hold.

2.1. RBF-based inner products

With the RBF regression, we seek to obtain an analytical expression of the velocity field $\mathbf{u}$ to approximate the inner products (2) and (3). We assume that training data is available at randomly scattered locations in space and uniform points in time $\{\mathbf{t}_k\}_{k=1}^{n_t}$. While this structure is not a strict requirement of the method, it reflects the typical format of data obtained from particle tracking velocimetry and particle-based numerical simulations. To construct the analytical expression, we represent the velocity field at time $\mathbf{t}_i$ as a linear combination of n_b RBFs:

$$\mathbf{u}(\mathbf{x}, \mathbf{t}_i) = \begin{pmatrix} u(\mathbf{x}, \mathbf{t}_i) \\ v(\mathbf{x}, \mathbf{t}_i) \end{pmatrix} \approx \begin{pmatrix} \gamma^T(\mathbf{x}; \mathbf{X}_c(\mathbf{t}_i), \mathbf{c}(\mathbf{t}_i)) \mathbf{w}_u(\mathbf{t}_i) \\ \gamma^T(\mathbf{x}; \mathbf{X}_c(\mathbf{t}_i), \mathbf{c}(\mathbf{t}_i)) \mathbf{w}_v(\mathbf{t}_i) \end{pmatrix}, \quad (6)$$

where $\gamma \in \mathbb{R}^{n_b}$ is a vector stacking the value of n_b radial basis functions $\{\gamma_m(\mathbf{x}; \mathbf{X}_{c,m}(\mathbf{t}_i), \mathbf{c}_m(\mathbf{t}_i))\}_{m=1}^{n_b}$, at an arbitrary location $\mathbf{x}$. The set of basis function is identified by the vector of collocation points $\mathbf{X}_{c,m}(\mathbf{t}_i)$ and the vector of shape parameters $\mathbf{c}_m(\mathbf{t}_i)$, while the vectors $\mathbf{w}_u, \mathbf{w}_v \in \mathbb{R}^{n_b}$ gathers the weights of each RBF for each velocity component. In this work, we use radial basis function of the form:

$$\gamma_m(\mathbf{x}; \mathbf{X}_{c,m}, \mathbf{c}_m) = \left(1 + \frac{\|\mathbf{x} - \mathbf{X}_{c,m}\|_2}{\mathbf{c}_m}\right)^5 \left(1 - \frac{\|\mathbf{x} - \mathbf{X}_{c,m}\|_2}{\mathbf{c}_m}\right)_+, \quad (7)$$

where $\|\bullet\|_2$ is the ℓ_2 norm of a vector, and the subscript $+$ denotes the truncated positive part of a function, i.e. $(a)_+ = a$ if $a > 0$ and $(a)_+ = 0$ if $a < 0$. For each snapshot, the weights can be computed via ridge regression [67, 68], or through physics-constrained regression [54, 58]. More generally, enhancements such as spectral filtering, boundary condition penalties, or divergence free for coupled velocity-pressure fields can also be embedded in this step [51, 52], with linear constraints that could be implemented using Lagrange multipliers [54, 57, 58]. For the test cases investigated in this work, which consist of densely sampled two-dimensional slices of inherently three-dimensional datasets, the improvement in accuracy obtained by enforcing constraints was not sufficient

to justify the associated increase in computational cost. We therefore adopt an unconstrained formulation here, while noting that such constraints can become highly beneficial in settings with lower spatial sampling density [54]. We use a compactly supported basis and solve the resulting least-squares problem with the iterative, sparse lsmr solver [69] available in the python library *scipy*.

The analytic approximation in (6) can be used together with Gauss-Legendre quadrature or Monte Carlo methods to compute the inner products in (3), as proposed in [63]. In this work, we follow the alternative approach of computing the integral directly from the RBF approximation, taking the simplifying condition of having time independent collocation points and shape functions, i.e. $\gamma(\mathbf{x}; \mathbf{X}_c(\mathbf{t}_i), \mathbf{c}(\mathbf{t}_i)) = \gamma(\mathbf{x}; \mathbf{X}_c, \mathbf{c})$. This assumption is valid if the domain boundaries are constant in time and if the data density is sufficiently uniform across the domain. Then, the inner product (3) between two velocity snapshots at $\mathbf{t}_i$ and $\mathbf{t}_j$ becomes

$$\begin{aligned} \langle \mathbf{u}(\mathbf{x}, \mathbf{t}_i), \mathbf{u}(\mathbf{x}, \mathbf{t}_j) \rangle_{\Omega} &\approx \langle \mathbf{u}(\mathbf{x}, \mathbf{t}_i), \mathbf{u}(\mathbf{x}, \mathbf{t}_j) \rangle_{\Omega,a} \\ &= \frac{1}{|\Omega|} \int_{\Omega} \sum_{k=u,v} \left(\sum_{m=1}^{n_b} \gamma_m(\mathbf{x}) \mathbf{w}_{l,m}(\mathbf{t}_i) \right) \left(\sum_{n=1}^{n_b} \gamma_n^*(\mathbf{x}) \mathbf{w}_{l,n}^*(\mathbf{t}_j) \right) d\Omega \\ &= \sum_{k=u,v} \sum_{m=1}^{n_b} \sum_{n=1}^{n_b} \mathbf{w}_{l,m}(\mathbf{t}_i) \mathbf{w}_{l,n}^*(\mathbf{t}_j) \underbrace{\left(\frac{1}{|\Omega|} \int_{\Omega} \gamma_m^*(\mathbf{x}) \gamma_n(\mathbf{x}) d\Omega \right)}_{:= \mathbf{I}_{\Omega}} \\ &= \sum_{k=u,v} \mathbf{w}_k^*(\mathbf{t}_i) \mathbf{I}_{\Omega} \mathbf{w}_k(\mathbf{t}_j), \end{aligned} \quad (8)$$

having defined the RBF based inner product $\langle \bullet, \bullet \rangle_{\Omega,a}$ and having introduced the integrating matrix $\mathbf{I}_{\Omega} \in \mathbb{R}^{n_b \times n_b}$ gathering the inner products of the basis functions.

The extension to inner products in time is analogous, relaxing the standard assumption of uniform temporal sampling and enabling decompositions that are continuous in both space and time. This work focuses on the discrete-time formulation, but we highlight that spatial RBF regression yields continuous spatial structures. The following sections describe the meshless implementation of the key decomposition methods.

2.2. Meshless algorithms

We introduce the meshless variant of the traditional POD (Sec. 2.2.1) and its spectrally constrained variants (Sec. 2.2.2), the Spectral POD (Sec. 2.2.3) and the Dynamic Mode Decomposition (Sec. 2.2.4). The extension of other grid-based decomposition methods to the RBF-induced meshless framework is conceptually straightforward and can be carried out following analogous steps.

2.2.1. Proper orthogonal decomposition (POD)

In the classic POD, the temporal structures of the modes are eigenfunction of the two point correlation function. The formulation for the continuous problem is usually presented for the so called “space-only” POD [1, 20, 70], but a similar derivation applies for the time domain.

The temporal structures $\psi_r(t)$ for the POD are taken to optimally represent the evolution of the field $\mathbf{u}(\mathbf{x}, t)$ at any location $\mathbf{x}$. To obtain these modes, we define the temporal correlation kernel as

$$K(t, t') = \int_{\Omega} \mathbf{u}(\mathbf{x}, t) \mathbf{u}(\mathbf{x}, t') d\mathbf{x}, \quad (9)$$

and the associated self-adjoint operator $\mathcal{K}$ acting on functions $\psi \in L^2([0, T])$ via:

$$(\mathcal{K}\psi)(t) = \int_0^T K(t, t') \psi(t') dt'. \quad (10)$$

With these definitions, the optimal temporal basis functions $\psi(t)$, which minimize the l_2 norm of the approximation error for any number of selected modes, are the solutions to the Fredholm integral eigenvalue problem:

$$\int_0^T K(t, t') \psi_r(t') dt' = \lambda_r \psi_r(t). \quad (11)$$

These bases are orthogonal under the continuous inner product (2). In the traditional setting of discrete time realizations available on a uniform grid $\{\mathbf{t}_k\}_{k=1}^{n_t}$, the continuous eigenvalue problem (11) is turned into a matrix eigenvalue problem for the time correlation matrix $\mathbf{K}_{i,j} = \langle \mathbf{u}(\mathbf{x}, \mathbf{t}_i), \mathbf{u}(\mathbf{x}, \mathbf{t}_j) \rangle_{\Omega, d}$. This matrix collects the inner product in space between all the available pairs of snapshots. Since $\mathbf{K} \in \mathbb{R}^{n_t \times n_t}$ is at least positive semidefinite by definition, the eigenvalue problem reads

$$\mathbf{K} = \mathbf{\Psi} \mathbf{\Lambda} \mathbf{\Psi}^T, \quad (12)$$

and the eigenvectors $\{\mathbf{\Psi}_r\}_{r=1}^{n_t}$ of $\mathbf{\Psi} \in \mathbb{R}^{n_t \times n_t}$ can be interpreted as discrete approximations, evaluated on the uniform time grid $\mathbf{t}$, of the continuous temporal eigenfunctions $\psi_r(t)$. These vectors are orthogonal under the inner product (4).

The proposed meshless variant of the POD replaces the inner product in the definition of $\mathbf{K}$ with the RBF-based inner product in (8), i.e. $\mathbf{K}_{i,j} \approx \langle \mathbf{u}(\mathbf{x}, \mathbf{t}_i), \mathbf{u}(\mathbf{x}, \mathbf{t}_j) \rangle_{\Omega, d}$, thus allowing to handle scattered data. In the simpler setting of fixed RBF bases the integrating matrix $\mathbf{I}_{\Omega}$ is independent of time, and the temporal correlation matrix can be computed as:

$$\mathbf{K} = \sum_{k=u,v} \mathbf{W}_k^T \mathbf{I}_{\Omega} \mathbf{W}_k, \quad (13)$$

where $\mathbf{W}_u, \mathbf{W}_v \in \mathbb{R}^{n_b \times n_t}$ collect the weights of all snapshots for each velocity component. Note that, with a slight abuse of notation, we use the same symbols $\mathbf{K}$ and $\mathbf{\Psi}$ to refer to the temporal correlation matrix and its eigenvectors, regardless of whether $\mathbf{K}$ is computed on grid data using the inner product (5) between snapshots, or via the RBF-based inner product between regressions. The distinction between the two is made clear in the following when strictly necessary.

We stress that the definition (13) also leads to a matrix that is at least positive semidefinite, hence having orthogonal eigenvectors under the inner product (4).

The second step in the traditional POD computation, following Sirovich's method [17], is the projection of the data onto the bases $[\mathbf{\Psi}_r]_{r=1}^{n_t}$ either using the continuous inner product $\langle \cdot \rangle_T$ in (2) for continuous data or the discrete one $\langle \cdot \rangle_{T,d}$ in (4) for gridded data. Since the orthogonality of the temporal structure $\mathbf{\Psi}_r$ is preserved by construction, the same approach remains valid for the meshless formulation proposed in this work. Therefore, it is possible to obtain an approximation of the continuous spatial structure for the u and v components associated to the r -th temporal structure, denoted as $(\phi_{u,r}(\mathbf{x})$ and $\phi_{v,r}(\mathbf{x})$ respectively, with the discrete inner product in time. For the velocity component u , for example, this reads:

$$\begin{aligned} \sigma_r \phi_{u,r}(\mathbf{x}) &= \langle u(\mathbf{x}, t), \mathbf{\Psi}_r(t) \rangle_T \\ &\approx \langle u(\mathbf{x}, \mathbf{t}_i), \mathbf{\Psi}_r(\mathbf{t}_i) \rangle_{T,d} \\ &\approx \frac{1}{n_t} \boldsymbol{\gamma}^T(\mathbf{x}) \sum_{i=1}^{n_t} \mathbf{w}_u(\mathbf{t}_i) \mathbf{\Psi}_r(\mathbf{t}_i) \\ &= \frac{1}{n_t} \boldsymbol{\gamma}^T(\mathbf{x}) (\mathbf{W}_u \mathbf{\Psi}_r) = \frac{1}{n_t} \boldsymbol{\gamma}^T(\mathbf{x}) \mathbf{w}_{\phi,u,r}, \end{aligned} \quad (14)$$

where σ_r is the r -th modal amplitude, $\mathbf{\Psi}_r(\mathbf{t}_i)$ is the i -th sample of the r -th temporal structure, $\boldsymbol{\gamma}^T(\mathbf{x})$ is the shortened notation for $\boldsymbol{\gamma}^T(\mathbf{x}; \mathbf{X}_c, \mathbf{c})$ and $\mathbf{w}_{\phi,u,r}$ is the weight vector for the RBF approximation of the u component of the r -th spatial structure $\boldsymbol{\phi}_r(\mathbf{x}) = [\phi_{u,r}(\mathbf{x}), \phi_{v,r}(\mathbf{x})]$.

We note that the amplitude σ_r serves as a normalization factor to ensure unitary norm of the spatial structures under a given inner product. A natural choice is to set

$$\sigma_r^2 = \left\langle \langle u(\mathbf{x}, t), \mathbf{\Psi}_r(t) \rangle_{T,d} \right\rangle_{\Omega}. \quad (15)$$

In the case of gridded data, where POD effectively reduces to an SVD of the snapshot matrix [14], one has $\sigma_r = \lambda_r^2$, with λ_r the eigenvalues in (12). This relationship arises from the symmetry of the inner product applied to both the column and row spaces of the snapshot matrix. In the proposed RBF-based formulation, however, this symmetry no longer holds: projections in time are defined under the inner product (4) while projection in space are defined under the inner product (8). Moreover, the temporal modes are orthogonal under (4) while the spatial ones are not orthogonal under (8), since:

$$\begin{aligned} \langle \boldsymbol{\phi}_r, \boldsymbol{\phi}_s \rangle_{\Omega, d} &\propto \int_{\Omega} \left(\sum_{l=1}^{n_b} w_{r,l} \gamma_l(\mathbf{x}) \right) \left(\sum_{l=1}^{n_b} w_{s,l} \gamma_l(\mathbf{x}) \right) d\Omega \\ &= \mathbf{w}_{\phi,r}^T \mathbf{I}_{\Omega} \mathbf{w}_{\phi,s}, \end{aligned} \quad (16)$$

where $\mathbf{w}_{\phi,i}$ and $\mathbf{w}_{\phi,j}$ are the RBF weights for the spatial structures following the time projection in (14). Thus we see that preserving the spatial orthogonality of the POD modes under (8) requires that the weight vectors of the RBF expansions of the spatial modes be orthogonal with respect to the inner product weighted by the integration matrix $\mathbf{I}_{\Omega}$. This condition could be enforced as an additional constraint during the RBF regression of all snapshots, using an iterative procedure that alternates between computing the temporal

correlation matrix $\mathbf{K}$, its eigenvectors Ψ , the associated RBF weights $\mathbf{w}_{\phi,i}$, and the inner products in (16). Nevertheless, since the loss of orthogonality was found to be minor, we leave iterative correction as a topic for future work. We note, however, that such an extension would become particularly relevant when constructing meshless methods for Galerkin-based reduced-order models.

2.2.2. Spectrally constrained PODs: the SPOD and mPOD

The Spectral POD by Sieber et al. [19] (hereafter referred to as *Sieber SPOD*) and the Multiscale POD by Mendez et al. [21] (hereafter *mPOD*) follow the same general steps as the traditional snapshot POD, with additional operations applied to the temporal correlation matrix $\mathbf{K}$. Specifically, Sieber SPOD applies a filter along the diagonals of $\mathbf{K}$ to promote a circulant structure, yielding eigenvectors that more closely resemble harmonic modes. The mPOD employs multiresolution analysis to decompose $\mathbf{K}$ into a sum of components at different temporal scales, each of which is equipped with its own POD. The reader is referred to the original articles for further details. The extension of these algorithms to the RBF-based meshless formalism introduces no additional steps beyond those already required for the traditional POD.

2.2.3. A harmonic POD: another SPOD

The SPOD of Towne et al. [20] (hereafter referred to as ‘Towne SPOD’) uses Welch’s method [71], and partitions the data in time into n_B different overlapping blocks of length n_f . A discrete Fourier transform (DFT) is carried out on each of these blocks. For the proposed meshless RBF based variance, the DFT of the data simply turns into a DFT of the weights if the RBF basis is fixed in time:

$$\hat{\mathbf{u}}(\mathbf{x}, \mathbf{f}_l) = \begin{pmatrix} \hat{u}(\mathbf{x}, \mathbf{f}_l) \\ \hat{v}(\mathbf{x}, \mathbf{f}_l) \end{pmatrix} \approx \begin{pmatrix} \gamma^\top(\mathbf{x}) \mathbf{W}_u \Psi_{\mathcal{F}} \\ \gamma^\top(\mathbf{x}) \mathbf{W}_v \Psi_{\mathcal{F}} \end{pmatrix} = \begin{pmatrix} \gamma^\top(\mathbf{x}) \hat{\mathbf{W}}_u \\ \gamma^\top(\mathbf{x}) \hat{\mathbf{W}}_v \end{pmatrix}, \quad (17)$$

where $\Psi_{\mathcal{F}} \in \mathbb{C}^{n_f \times n_f}$ is the Discrete Fourier Transform matrix [72], defined as $\Psi_{\mathcal{F},l}(\mathbf{t}_k) = \exp(2\pi i \mathbf{f}_l \mathbf{t}_k) / \sqrt{n_f}$ with $\{\mathbf{f}_l = l/n_f \mathbf{f}_s\}_{l=1}^{n_f}$ the set of uniformly spaced frequencies and $\hat{\mathbf{W}}_u$ and $\hat{\mathbf{W}}_v$ the row-wise Fourier transforms of the weight matrices. For every (discrete) frequency in the set $\{\mathbf{f}_l\}_{l=1}^{n_f}$, the DFTs of each block are then stacked together. For the velocity u , for example, this results in:

$$\begin{aligned} \hat{\mathbf{u}}_{\mathcal{W}}(\mathbf{x}, \mathbf{f}_l) &= [\hat{u}^{(1)}(\mathbf{x}, \mathbf{f}_l), \dots, \hat{u}^{(n_B)}(\mathbf{x}, \mathbf{f}_l)] \\ &\approx \gamma^\top(\mathbf{x}) [\hat{\mathbf{w}}_u^{(1)}(\mathbf{f}_l), \dots, \hat{\mathbf{w}}_u^{(n_B)}(\mathbf{f}_l)] \\ &= \gamma^\top(\mathbf{x}) \hat{\mathbf{W}}_{\mathcal{W},u}(\mathbf{f}_l), \end{aligned} \quad (18)$$

where $\hat{\mathbf{w}}_u^{(p)}(\mathbf{f}_l)$ collects the DFT of the p -th Welch block at the l -th frequency and $\hat{\mathbf{W}}_{\mathcal{W},u} \in \mathbb{C}^{n_b \times n_B}$ gathers the individual Welch contributions at frequency $\mathbf{f}_l$ of the Fourier transform of the weights, as in the original work [20].

This Welch block is formed for both velocity components and used to form a correlation matrix $\mathbf{M} \in \mathbb{C}^{n_B \times n_B}$ with an inner

product in space:

$$\begin{aligned} \mathbf{M}(\mathbf{f}_l) &= \langle \hat{\mathbf{u}}(\mathbf{x}, \mathbf{f}_l), \hat{\mathbf{u}}(\mathbf{x}, \mathbf{f}_l) \rangle_{\Omega,a} \\ &= \sum_{k=u,v} \hat{\mathbf{W}}_{\mathcal{W},k}^*(\mathbf{f}_l) \mathbf{I}_\Omega \hat{\mathbf{W}}_{\mathcal{W},k}(\mathbf{f}_l), \end{aligned} \quad (19)$$

where the derivation from the inner product to the integration matrix $\mathbf{I}_\Omega$ follows the same steps as (8).

As in the original work, an eigenvalue decomposition is then carried out on this matrix for each frequency, that is $\mathbf{M}(\mathbf{f}_l) = \mathbf{\Theta}(\mathbf{f}_l) \Sigma^2(\mathbf{f}_l) \mathbf{\Theta}^*(\mathbf{f}_l)$, and the r -th spatial mode of u associated to frequency $\mathbf{f}_l$ is obtained via

$$\phi_{u,r}(\mathbf{x}, \mathbf{f}_l) = \frac{1}{\sigma_r(\mathbf{f}_l)} \gamma^\top(\mathbf{x}) \hat{\mathbf{W}}_{\mathcal{W},u}(\mathbf{f}_l) \mathbf{\Theta}_r(\mathbf{f}_l), \quad (20)$$

where $\mathbf{\Theta}_r(\mathbf{f}_l)$ is the r -th column of $\mathbf{\Theta}(\mathbf{f}_l)$. The result is a total of n_B modal amplitudes and spatial structures at each frequency $\mathbf{f}_l$.

2.2.4. Dynamic Mode Decomposition (DMD)

Many algorithms for computing the DMD have been proposed [73]. We here consider the POD-based formulation proposed by Schmid [18] which seeks to identify the eigenvalue decomposition of a reduced propagator $\tilde{\mathbf{S}}$ that best approximates, in a least square sense, the dynamics in the POD-reduced space using a linear system. In the traditional grid based formulation, denoting as $\mathbf{D}_1$ and $\mathbf{D}_2 \in \mathbb{R}^{n_s \times (n_t-1)}$ the snapshot matrices containing the snapshots from 1 to $n_t - 1$ and from 2 to n_t and as $\mathbf{P} \in \mathbb{R}^{n_s \times n_s}$ the (unknown) full propagator such that $\mathbf{D}_2 = \mathbf{P} \mathbf{D}_1$, the reduced propagator is defined as $\tilde{\mathbf{S}} = \tilde{\mathbf{\Phi}}_{\mathcal{P}}^\top \mathbf{P} \tilde{\mathbf{\Phi}}_{\mathcal{P}}$, where $\sim$ denotes quantities computed from a reduced number of modes. Using the truncated POD (SVD) to compute the pseudoinverse $\mathbf{D}^{\dagger+} = \tilde{\mathbf{\Psi}}_{\mathcal{P}} \tilde{\mathbf{\Sigma}}_{\mathcal{P}} \tilde{\mathbf{\Phi}}_{\mathcal{P}}^{-1}$, and introducing it in the least square definition of the propagator $\mathbf{P} = \mathbf{D}_2 \tilde{\mathbf{\Psi}}_{\mathcal{P}} \tilde{\mathbf{\Sigma}}_{\mathcal{P}} \tilde{\mathbf{\Phi}}_{\mathcal{P}}^{-1}$, the reduced propagator becomes [18, 73]:

$$\tilde{\mathbf{S}} = \tilde{\mathbf{\Phi}}_{\mathcal{P}}^\top \mathbf{D}_2 \tilde{\mathbf{\Psi}}_{\mathcal{P}} \tilde{\mathbf{\Sigma}}_{\mathcal{P}}^{-1}, \quad (21)$$

The meshless variant of (21) can be easily derived by replacing $\tilde{\mathbf{\Sigma}}_{\mathcal{P}}$ and $\tilde{\mathbf{\Psi}}_{\mathcal{P}}$ with their meshless counterpart and replacing the inner product $\tilde{\mathbf{\Phi}}_{\mathcal{P}}^\top \mathbf{D}_1$ with the RBF-based one.

A simpler and computationally cheaper alternative is offered by variants of this idea that were developed earlier in atmospheric science, most notably the Principal Oscillation Pattern (POP) and Principal Interaction Pattern (PIP) frameworks [74–76], which also estimate a linear propagator directly from time-lagged snapshots.

These allow to compute the reduced propagator without involving inner product in the space domain. The underlying idea is to assume that the two shifted snapshot sequences (the snapshot matrices $\mathbf{D}_1, \mathbf{D}_2$ in traditional DMD, or their space continuous formulation in the meshless approach) share the same POD and only differ in the temporal structures. Then, it can be shown (see [77]) that (21) becomes:

$$\tilde{\mathbf{S}} = \tilde{\mathbf{\Sigma}}_{\mathcal{P}}^{-1} \tilde{\mathbf{\Psi}}_{\mathcal{P},1}^\top \tilde{\mathbf{\Psi}}_{\mathcal{P},2} \tilde{\mathbf{\Sigma}}_{\mathcal{P}}, \quad (22)$$

where $\mathbf{\Psi}_{\mathcal{P},1} = \{\mathbf{\Psi}_{\mathcal{P}}(\mathbf{t}_i)\}_{i=1}^{n_t-1}$ and $\mathbf{\Psi}_{\mathcal{P},2} = \{\mathbf{\Psi}_{\mathcal{P}}(\mathbf{t}_i)\}_{i=2}^{n_t}$, with $\mathbf{\Psi}_{\mathcal{P}}$ the temporal structure of the data POD and $\tilde{\mathbf{\Sigma}}_{\mathcal{P}}$ the associated




Figure 1: Test case 1. Mean subtracted velocity magnitude for the impinging jet. Example PIV snapshot (top) and the time step interpolated onto scattered points with downsampling factor $k = 1$ (middle) and $k = 16$ (bottom).

amplitude matrix. The meshless variant of (22) follows directly from the meshless POD construction and requires no additional RBF inner products beyond the single POD computation.

Finally, as for the grid-based DMD, the eigenvalue decomposition of the reduced propagator is carried out $\tilde{S} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^*$ and the resulting spatial structures are computed from spatial projection of the POD structure. For the meshless DMD, this projection is computed from the RBF weights, the POD of the data, and the eigenvectors of the reduced propagator. For the u-component of the spatial structure, for example, one has:

$$\phi_{u,r}(\mathbf{x}) = \tilde{\phi}_p(\mathbf{x}) \mathbf{q}_r = \frac{1}{\sigma_r n_t} \gamma^T(\mathbf{x}) \mathbf{W}_u \Psi_{p,r} \mathbf{V}_r. \quad (23)$$

3. Selected test cases

The meshless decompositions are analyzed using two different datasets. The first dataset contains experimental data from PIV and features two typical experimental challenges: The spatial resolution is coarse because of the low-pass filtering of the interrogation windows, and the range of relevant frequencies in time is close to the Nyquist limit. This is common for experiments where increasing the spatial or temporal resolution arbitrarily is limited by hardware. The second dataset contains numerical data from large eddy

simulations (LES) and features the contrary: The data are temporally well resolved to capture relevant temporal frequencies and contains spatial scales down to the filtering frequency of the LES model.

For both datasets, we obtain the scattered datasets by interpolating the original (gridded) data. We interpolate the data with different downsampling factors k , where $n_p = n_g/k$, simulating different data densities. Before interpolating, we remove the mean of the gridded data, since both datasets are stationary. If the gridded mean is unavailable, it can be subtracted from the data ensemble through binning [78], or another RBF regression [58].

3.1. Test case 1: Impinging jet from particle image velocimetry

The first test case uses experimental data from PIV of a jet impinging on a plate, denoted as ‘Jet’ in the remainder of the manuscript. The jet flow features three distinct flow patterns: the nozzle exit, the impinging region close to the wall, and the developing wall jet. Each of these occurs at a different spatial location, with a different frequency in time and in space. The measurements were conducted at the von Karman Institute and featured a rectangular nozzle at a distance of $Z = 30$ mm with a size of $H \times W = 4 \times 250$ mm and an exit velocity of $U_0 = 6.5$ m/s, corresponding to a Reynolds number of 1733. Further details can be found in Mendez et al. [21]. The domain of interest had a size of 20×38 mm², and 2000 snapshots were acquired at a frequency of 2 kHz. This corresponds to a dimensionless time $t U_0/H$ in the range $[0, 1625]$ with a temporal resolution $\Delta t U_0/H = 0.81$. Relevant coherent structures are present up to a dimensionless frequency of approximately $St = 0.4$, with St the Strouhal number, while the Nyquist frequency is $St = fH/U_0 = 0.61$.

The resulting velocity is available in $n_g = 6840$ points on a regular grid of 60×114 in z and x -direction with a vector pitch of $\Delta x = 0.333$ mm. Since the grid is regular, we use piecewise cubic Hermite interpolating polynomials to produce scattered data. Three different datasets are generated with different downsampling factors $k \in [1, 4, 16]$. The first panel of Figure 1 shows the (mean-subtracted) magnitude of a PIV snapshot, illustrating the different turbulent length scales in the free shear layer and the developing wall jet. The bottom two subpanels show the same snapshot interpolated with $k = 1$ and $k = 16$. For the latter, the different vortices are hardly visible by eye.

3.2. Test case 2: Transitional airfoil from large eddy simulations

The second test case uses data from an LES of the flow past a NACA 0012 at 6° angle of attack at a Reynolds number of 23 000, denoted as ‘Airfoil’ in the remainder. In these conditions, the laminar flow separates near the leading edge, transitions to turbulence, and reattaches close to the trailing edge. The Kelvin-Helmholtz instability triggers vortex shedding over the separation bubble, leading to a wide spectrum over the airfoil and in the wake. Further details can be found in Yeh and Taira [79], Towne et al. [80]. From the




Figure 2: Test case 2. Mean subtracted velocity magnitude for the airfoil flow. Example LES snapshot (top) and the time step interpolated onto scattered points with downsampling factor $k = 2$ (middle) and $k = 32$ (bottom). The aspect ratio is slightly increased for better visualization.

full simulation environment, we extract data in a domain $(x/c, y/c) \in [-0.1, 2] \times [-0.2, 0.2]$, where c is the chord of the airfoil. The data are available in 16 000 snapshots spanning a dimensionless time of $t c/U_\infty \in [0, 165]$ corresponding to a dimensionless timestep of $\Delta t c/U_\infty = 0.0104$. Relevant structures appear up to a frequency of $St \approx 4$, giving this test case a much higher temporal resolution with a Nyquist frequency of $St = 48$.

We sample span-wise averaged velocities u and v in the $n_g = 60\,844$ unstructured LES points. Since this grid is irregular, we resort to an RBF interpolation with a linear kernel to avoid ill-conditioning caused by the hole in the domain. To avoid excessive oversampling in regions with large grid spacing (i.e., far from the airfoil), we interpolate this data with the three downsampling factors $k \in [2, 8, 32]$. The first panel of Figure 2 shows the (mean subtracted) velocity magnitude, which displays the characteristics of the laminar separation bubble over the chord and the resulting turbulent wake. The lower panels of the same figure show the same time step, interpolated with $k = 2$ and 32 . For the dense case, the turbulent structures are again clearly visible over the airfoil and in the wake. For $k = 32$, only 1900 points are available, which makes it almost impossible to spot structures in the wake or over the latter half of the chord.

4. Results

We compare the RBF-based modal decompositions with the reference data and binning. For the jet flow, the gridded PIV is taken as the ground truth, acknowledging the data being smoothed and contaminated by noise. The spatial structures are visualized in the grid points of the underlying data, meaning the RBFs are evaluated in these points and binning uses them




Figure 3: Normalized L_2 error in the temporal correlation matrix $\mathbf{K}$ (24) versus the downsampling factor for binning and RBFs for the jet (left) and the airfoil (right). The dashed circles correspond to the downsampling factors analyzed in detail in the remainder of the article.

as centroids. The bin size is chosen such that each bin contains 10 particles on average; a trade-off between accuracy and noise filtering, used in other benchmarks [53]. Finally, the sample-based inner products of binning and the ground truth of the airfoil are computed with a weighted inner product based on the cell volume [14].

For both datasets, the RBF regression uses the same parameters. We choose a ratio between training points and RBFs of $n_p/n_b = 1.5$ and select a uniform shape factor c such that an average of 40 data points is underneath each RBF. They are placed semi-randomly in space with Halton points, and collocation points are removed within the airfoil to mitigate ill-conditioning. The ridge regression penalty is set to correspond to $\alpha_{\text{ridge}} = 0.01$ for a direct least squares solver [67, 68]. These hyperparameters are selected based on similar regressions (e.g. [51, 56]), but we observed the results to depend relatively little on the hyperparameters. A reasonable regression always resulted in reasonable decompositions.

For the Towne SPOD, we use Welch blocks with 50 % overlap with $N_f = 200$ and 2000 samples for the jet and the airfoil. The results likewise varied little for any reasonable range of overlaps. For the Sieber SPOD and mPOD, we respectively tested different filter cut-offs and filter banks, and we show results sparingly where appropriate. Finally, we compute the DMD with 200 and 2000 POD modes for the jet and the airfoil.

The results are presented not split by the different methods, but by their common steps. That is, we start with the correlation matrix based on the spatial inner product and then continue with the decomposition of this matrix into modal amplitudes and the temporal bases. From there, we obtain the spatial structures, either by projecting the data onto the temporal basis structures (POD/mPOD/SPOD) or by computing the reduced propagator and its eigendecomposition (DMD).

4.1. The correlation matrix

As a first estimate of how each method performs, we analyze the normalized error in the temporal correlation matrix $\mathbf{K}$, taken




Figure 4: Temporal correlation matrix $\mathbf{K}$ of the jet data (left) and airfoil data (right). The upper left panel on each side shows the gridded (ground truth) data while the second and third row respectively show the binning and RBF results. The columns on each side show different downsampling factors, and the colorscale is shared across all subpanels. For better visualization, only the first 200 time steps of each dataset are displayed.

as:

$$\|\mathbf{K}\|_{\text{err}} = \frac{\|\mathbf{K}_{\text{exact}} - \mathbf{K}\|_2}{\|\mathbf{K}_{\text{exact}}\|_2}, \quad (24)$$

where $\mathbf{K}$ is the correlation matrix from either binning or the RBF-based inner product and the subscript 'exact' refers to the gridded data from the PIV and LES.

These errors over the downsampling factor k are visualized in Figure 3 for the jet and airfoil on the left- and right-hand side. The RBF error is consistently smaller than the binning one, reaching values below 6 % for the jet and below 1 % for the airfoil. The gap between the two methods is almost constant for larger k , but as the downsampling decreases, the gap narrows for the jet while it widens for the airfoil. We attribute this to the nature of the data: The PIV data are contaminated and biased by reflections near the wall. The RBFs partially filter these errors and thus never fully match the correlation matrix of the PIV. In contrast, the LES is noise-free and resolves more scales which is why more data also results in higher accuracy for the RBFs. The same is true for binning, although the benefits at lower k are smaller.

The dashed circles in Figure 3 indicate the three downsampling factors mentioned in section 3, and we analyze only these in the following to reduce the number of subpanels in the subsequent figures. For these three conditions, Figure 4 shows the correlation matrix zoomed in on the first 200 timesteps. The left-hand side shows the jet and the right-hand side the airfoil results, while the rows respectively correspond to the gridded, binning, and RBF data. The title of the columns indicates the different downsampling factors. From the gridded

results, the difference between the datasets is clearly visible. The jet has a coarse temporal resolution, resulting in small “structures” in the correlation matrix. In contrast, the airfoil has a fine temporal resolution, resulting in a visually much smoother matrix in the focused region. If the matrix is viewed with more timesteps, the airfoil case likewise shows a similarly “complex” pattern as the jet, albeit at much higher resolution.

For the smallest downsampling factor, both binning and the RBFs reproduce the respective patterns of the gridded data well. Then, as the downsampling factor increases, we observe two things. First, the magnitude of the correlation matrix decreases, which we attribute to the spatial low-pass filtering of both methods. Large fluctuations in space are either not present in the data or averaged out, leading to a smaller velocity magnitude and thus, a dampened correlation matrix. Second, the fine-scale structures in the matrix are smoothed out. This happens most severely for binning of the jet due to the coarse temporal resolution. This apparent smoothing is stronger for binning since it is a straightforward low-pass filter, while the regression properties of the RBFs mitigate spatial smoothing to some extent. These results also agree with the error curves in Figure 3, which are linked to a decreasing magnitude and thus, a larger error.

The correlation matrices of the Towne SPOD show similar patterns, which is why they are omitted. The RBFs also consistently reach a lower error in the magnitude of the correlation matrices across all frequencies, which is again linked to the low-pass filtering and attenuation of binning. Low-pass filtering in space evidently dampens the Fourier amplitudes, and hence the correlation matrix $\mathbf{M}$.




Figure 5: Diagonal of the power spectral density of the temporal correlation matrix $\mathbf{K}$ over the Strouhal number for the jet (top) and the airfoil (bottom). The power spectral density is computed using Welch’s method with Hanning windows, 75 % overlap and 500 (Jet) and 4000 (Airfoil) samples per segment. The columns correspond to increasing downsampling factors according to the title of each subpanel. The legend is shared across all subpanels and in regions of strong overlap, the curves match almost completely.

To assess the meshless mPOD and Sieber SPOD, we compare the spectral content of the correlation matrix $\hat{\mathbf{K}}$ by computing its power spectral density (PSD). Figure 5 displays the diagonal of this matrix for the jet (top) and the airfoil (bottom). For visualization purposes, we use Welch’s method with Hanning windows, 75 % overlap, and 500 (Jet) and 4000 (Airfoil) samples per segment; the PSD without windowing shows the same trend.

For the jet, both methods yield similar results for $k = 1$. The curves overlap almost completely and are indistinguishable except for the highest frequencies close to the Nyquist limit at $St = 0.5$. The wide spectrum is clearly visible, with pronounced peaks up to $St = 0.3$. For $k = 4$, binning yields a slightly attenuated spectrum, particularly for frequencies above $St = 0.1$. This attenuation is further amplified for $k = 16$: The entire spectrum is dampened, the peak around $St = 0.2$ is barely visible, and the footprint of the Kelvin-Helmholtz rollers in the shear layer at $St = 0.3$ is not visible. In contrast, the RBFs yield an almost perfect spectrum for $k = 4$ and show only minor differences for $k = 16$. Significant peaks are still present at all frequencies, and all relevant features are recovered.

We expect these observations to also affect the mPOD and Sieber SPOD. Binning cannot recover features at high Strouhal

numbers, and we can hardly expect any filtering operation on the correlation matrix to do so. The same happens for the mPOD: what is lost in space and time cannot be recovered, no matter the decomposition. Since the RBFs yield a better spectrum of $\mathbf{K}$, we expect the meshless mPOD and Sieber SPOD to work equally well as in the gridded case.

In contrast, the airfoil case shows only minor differences in the spectrum in the relevant frequencies, even for $k = 32$. A peak at $St \approx 1.47$ identifies the frequency of the laminar separation bubble [3], and the harmonic peaks up to $St = 3.5$ are clearly visible across all cases. Higher frequencies tell a different story. For both methods, the spectrum flattens and resembles white noise. For binning, this transition toward white noise happens at lower frequencies and is preceded by a dampened amplitude. Conversely, the RBFs follow the grid data almost perfectly until they drop off to white noise, the earliest frequency being $St = 10$ for $k = 32$. We link this different behavior of the RBFs to the different datasets: Compared to the jet, the amplitude range of the PSD is twice as large. For the jet, almost all frequencies are recovered well since they contain significant (spatial) features, such as vortices, while the high frequencies of the airfoil show no significant features and have small amplitude. Because of this, we expect




Figure 6: Modal amplitudes from the POD over the mode number r for the jet (top) and the airfoil (bottom). The columns correspond to increasing downsampling factors according to the title of each subpanel. Only the first 200 and 2000 modes are shown for the jet and the airfoil, respectively. The legend is shared across all subpanels.

both the mPOD and the Sieber SPOD to work well for the dominant features of both binning and RBF.

4.2. The modal amplitudes

Any error in the correlation matrices $\mathbf{K}$ and $\mathbf{M}$ will also influence their eigendecomposition. Since the eigenvectors are orthonormal, we expect a different magnitude of the matrix to affect only the modal amplitudes, whereas a different structure of the matrix can affect both the amplitudes and the eigenvectors. To investigate the first hypothesis, Figure 6 shows the modal amplitudes of the eigendecomposition of $\mathbf{K}$ for the jet (top row) and the airfoil (bottom row), while the columns correspond to different downsampling factors according to the title of each subpanel. For better comparison, we only show the first 200 and 2000 modes for the jet and airfoil.

For the jet, both binning and RBF match the result of the gridded POD well for $k = 1$. Modes above 100 are slightly attenuated, but the curves show the same trends. As the downsampling increases, the modal amplitudes are dampened further, albeit stronger for binning than for RBF. For $k = 16$, already the third mode of binning has only 75 % of the grid amplitude, and this gap is almost constant over a wide range of modes. In contrast, the RBFs show similar dampening only for modes above 100, albeit with a steeper drop-off than binning. The airfoil curves show similar trends, although less pronounced. For $k = 2$, slight attenuation occurs only above mode 1000, and even at $k = 32$, binning captures the first 30 modes well. As before, the RBFs match the grid curve

better, but then have a sudden, steep drop in the amplitude after approximately 500 modes for $k = 32$. These differences between binning and RBF confirm the hypothesis about the correlation matrix: A dampened correlation matrix leads to dampened modal amplitudes.

One consistent issue is the steep drop-off in the modal amplitudes of the RBFs at high mode numbers. Larger eigenvalues beyond the range shown here are multiple orders of magnitude smaller. We attribute this to the numerical stability of the solver since a Sieber SPOD with a cut-off frequency at $St = 25$ (airfoil) already removes this sharp drop-off. The mPOD shows similar results. Any range of reasonable parameters results in modal amplitudes matching the gridded result much better than binning. The other observations remain, which is why we omit to show the modal amplitudes of these POD variants.

The decomposition of the correlation matrix $\mathbf{M}$ for the Towne SPOD yields multiple modal amplitudes at each frequency. Figure 7 depicts the two largest modal amplitudes over the Strouhal number, where the first and second rows correspond to the jet and the third and fourth to the airfoil. The columns again indicate the downsampling factor k , increasing from left to right.

For the jet test case, we obtain similar results as for the spectral analysis of the temporal correlation matrix. At small k , both methods yield more or less the same spectrum for the first and second modes. Then, as k increases, binning attenuates the spectrum. For $k = 4$, the differences are minor, but for $k = 16$,




Figure 7: First and second modal amplitude of the Towne SPOD over the Strouhal number for the jet (row 1 and 2) and the airfoil (row 3 and 4). The columns correspond to increasing downsampling factors according to the title of each subpanel. The range of frequencies matches Fig. 5, and the legend is shared across all subpanels.

the spectrum is flattened severely at high Strouhal numbers. However, compared to the spectral analysis of $\mathbf{K}$ (c.f. figure 5), a small peak at $St = 0.3$ can still be identified for binning. This difference is expected, since the Towne SPOD is a more sophisticated analysis than a simple PSD of $\mathbf{K}$. Conversely, the RBFs almost perfectly recover the whole spectrum, with only some slight differences at high frequencies beyond $St = 0.4$. This behavior is consistent across all mode numbers: Binning gives dampened spectra while the RBFs match almost perfectly.

The airfoil results agree with the previous analysis. In both the first and the second modes, the modal amplitudes of binning are slightly dampened. The higher k , the earlier this attenuation happens: at $St = 3$ in the sparsest case. Afterward, the spectrum flattens to white noise with the cut-off frequency again decreasing with increasing k . In comparison, the RBFs either completely follow the spectrum ($k = 1$) or follow it up to a cut-off frequency and then transition with a sharp kink toward

white noise. This cut-off frequency is above $St = 10$ for both modes and always higher for RBF than for binning. Moreover, it is almost constant for all modes. In summary, the SPOD results agree well with the spectral analysis of the temporal correlation matrix: The RBFs follow the spectrum of the grid more faithfully and up to a higher Strouhal number.

4.3. The temporal basis

For the POD, mPOD, and Sieber SPOD, the eigendecomposition of $\mathbf{K}$ gives the temporal structures. Figure 8 shows the PSD of the third POD mode for the jet (top) and the airfoil (bottom), while the columns again correspond to an increasing k from left to right, as indicated by the titles. For better comparison, the spectra are again visualized with the same Welch settings as in Figure 5.

The chosen mode of the jet contains a pronounced peak around $St = 0.15$, and both methods recover it equally




Figure 8: Power spectral density of the third POD mode over the Strouhal number for the jet (top) and the airfoil (bottom). The power spectral density is computed using Welch’s method with Hanning windows, 75 % overlap and 500 (Jet) and 4000 (Airfoil) samples per segment. The columns correspond to increasing downsampling factors according to the title of each subpanel. The legend is shared across all subpanels and in regions of strong overlap, the curves match almost completely.

well. Compared to the Nyquist frequency, this frequency is sufficiently small such that both methods recover it well for $k = 1$. However, as the downsampling increases, binning shows spectral mixing: the low-frequency components have a higher amplitude than the gridded and RBF versions. For $k = 4$, the peak at $St = 0.15$ can still be identified, but for $k = 16$, it is completely gone. The RBFs also show signs of spectral mixing at $k = 16$, with a slightly increased low and high-frequency component, but not as severe as binning. The third mode still clearly matches the gridded case.

The temporal structures of the mPOD and Sieber SPOD confirm the observations of the temporal correlation matrix, which is why they are omitted. At high k , high frequencies are diminished, which also translates into both POD variants. At $k = 16$, no filter for the SPOD and neither the mPOD can recover the features of the Kelvin Helmholtz instability sufficiently well. What is lost cannot be recovered, no matter how sophisticated the decomposition is.

For the airfoil, spectral mixing is not present in the low-frequency content of the first modes. The peak at $St \approx 1.5$ and the general shape of the spectrum are recovered almost perfectly for both methods and all downsampling factors. The only issues arise at high frequencies above $St = 10$. The spectra

of RBF and binning are again white noise, with a decreasing cut-off frequency for increasing k . We observe this white noise band in all modes; the cut-off frequency slightly varies but always has the same order of magnitude. Compared to the white noise behavior of $\mathbf{K}$ and the Towne SPOD (c.f. Figs. 5 and 7), the cut-off frequency is smaller. The RBFs still perform slightly better than binning, but the differences are minor: significant features are located low small frequencies compared to the Nyquist limit because of the excellent temporal sampling. This POD mode, and low POD modes in general, are identified well.

However, this broadband, white noise behavior only affects the classical POD since it is not constrained or penalized in the frequency content of each mode. For the mPOD, high frequencies are filtered naturally due to the different filter banks; modes only contain white noise if they are from the filter band above $St = 10$. Yet, since no relevant features are present at this frequency, the results of the mPOD are barely influenced by this. Neither are the results of the Sieber SPOD. For different finite impulse response filters, we observed the same trend: If the white noise band is present after a certain frequency, a filter with a cut-off frequency at this frequency always removes the band without affecting the dominant features of the mode. However these are minor details overall. The temporal resolution

of the airfoil dataset is sufficiently large so that no relevant feature is drowned noise. In the jet, this behavior is not visible since almost all frequencies have relevant content, and the PSD spans fewer orders of magnitude.

4.4. The spatial basis

We close the discussion of the different PODs by analyzing the resulting spatial structures. To this end, Figure 9 shows the third spatial structures, corresponding to the temporal mode of Figure 8. The upper and lower part of the figure show the modes of the jet and the airfoil with the top left subpanel showing the gridded ground truth and the second and third row the results from binning and the RBFs. The columns again correspond to different downsampling factors k .

For the jet, at $k = 1$, both binning and RBF show excellent agreement with gridded data. The Kelvin-Helmholtz rollers of the jet are clearly visible and look almost the same for all three methods. When fewer data points are available at $k = 4$, the RBFs remain largely unchanged, while we can observe two things for binning. First, small structures of the developing wall jet at the bottom become visible. This agrees well with the observations in Figure 8, where spectral mixing produced a temporal structure with multiple scales in time, leading to a spatial structure with multiple scales in space. The spectral mixing also likely causes the second observation: the Kelvin-Helmholtz rollers are out of phase compared to before. Space-only PODs always need at least two modes in quadrature to represent vortex shedding, but the respective spatial and temporal frequencies can also be present over a broader range of frequencies. The classic POD is affected most by this, while the slightly constrained Sieber and mPOD can somewhat mitigate this effect. Finally, at $k = 16$, the Kelvin-Helmholtz instability is barely visible, and the mode is dominated by the wall jet for binning, again agreeing with the temporal structures from before. Intermediate downsampling factors between the three displayed ones show a gradual shift between these two extremes for the mode. In contrast, the RBFs only show an additional minor near-wall structure—the Kelvin-Helmholtz rollers are still clearly visible with unchanged frequency or amplitude. In conclusion, if the temporal structures do not match, we can hardly expect the spatial structures to do so since they are computed from the projection of the data onto the temporal modes.

Conversely, for the airfoil, the temporal structures match well for all k and both methods, save for some broadband content in the spectrum. The spatial structures in the lower half of Figure 9 clearly all show similar vortex shedding in the wake and a dominant structure over the latter half of the chord. While this is a qualitatively good agreement, this test case shows a different error of any of the PODs. The spatial basis is computed by projecting the data (or weights) onto the temporal structures. Any low-pass filtering in the data thus also affects the spatial structures, even if the temporal structures and modal amplitudes match perfectly. This is clearly visible in the fine-scale structures between the leading edge and half the chord length: a thin structure is visible for the gridded case, and also binning and RBF at $k = 2$. This is associated with the edge

of the laminar separation bubble, and the resulting structures in the latter half of the chord are associated with the transitioning shear layer. With increasing k , these small structures start to be smoothed: at $k = 32$, the structures of separation are gone, and the ones of the transitioning shear layer are attenuated. In contrast, the structures for the RBFs only have a slightly diminished amplitude and are still clearly visible. We found this behavior to be consistent across all test cases and all variants of the POD: If the correlation matrices $\mathbf{K}$ or $\mathbf{M}$ and their eigendecomposition match the gridded case, then so do their spatial structures, save for some spatial smoothing depending on the data mapping method. Since the RBFs are more sophisticated and can better fill gaps than binning, they yield more accurate spatial structures.

In summary, the results of the different PODs clearly show the limitations of both binning and RBF. If the spectrum of the data contains significant features up to the acquisition frequency, as it does for the jet, good spatial mapping is required. If the resolution is too coarse, binning exhibits increased spectral mixing and thus produces poor spatial structures. The RBFs are more sophisticated and can thus recover the relevant spatial and temporal structures, even in challenging conditions. If the relevant frequencies are significantly smaller than the Nyquist limit, we can afford a simpler spatial mapping since the high temporal resolution helps to recover the correct temporal and spatial modes. However, even if the modal amplitudes and temporal structures are perfect, the spatial structures can still be oversmoothed since they are computed by projecting in space. Therefore, RBFs remain superior to binning since they result in more accurate spatial structures across all data densities.

4.5. The Dynamic Mode Decomposition

The DMD uses the temporal structures and modal amplitudes from the POD to compute the reduced propagator. Errors in these two quantities will thus also affect the DMD. Figure 10 shows the resulting eigenvalues λ of the propagator in the complex plane, for the jet on the left and for the airfoil on the right. The top left panel shows the grid result, while the second and third rows again correspond to binning and RBF, with the columns indicating different downsampling factors k . The red dot in each figure corresponds to a frequency whose spatial structure is analyzed further below.

For the jet at low k , both methods recover well the distribution of modes from the gridded data. Since the DMD is computed through the POD, only low-frequency modes with large real parts lie close to the unit circle. High-frequency modes are severely dampened. As k increases, the amplitude of most eigenvalues decreases, and only few modes remain pronounced and dominant. This occurs more strongly for binning where at $k = 16$, most eigenvalues are close to the origin and thus contribute only marginally to the decomposition. For the RBFs, more modes lie closer to the unit circle, and their magnitude is also larger for larger values of k . These observations agree well with the modal amplitudes in Figure 6. For the jet, the higher modes at $k = 16$ quickly have an attenuated amplitude for binning, while the ones of




Figure 9: Spatial structures of the third POD mode for the jet (top part) and the airfoil (bottom part). The top left subpanel each shows the ground truth from the gridded data and the rows correspond to the method labeled on the left. The columns correspond to increasing downsampling factors according to the title of each subpanel and the range of contours is shared across the panels of each case. For the airfoil, the aspect ratio is slightly increased for better visualization.

the RBF remain close to the gridded ones up to mode number 100. These dampened amplitudes evidently also affect the DMD eigenvalues. They are most dominant at low frequencies, corresponding well to the spectra of the correlation matrix $\mathbf{K}$ in Figure 5, where high frequencies are strongly dampened for binning but not for the RBFs. These properties are inherited via the POD of $\mathbf{K}$ to the eigenvalues of the DMD propagator.

The airfoil observations are similar. For the gridded case, more modes are close to the unit circle, even modes with a high Strouhal number. For $k = 2$, binning already has many modes with smaller magnitudes, while the RBFs reproduce the apparent pattern well. As k increases, many modes are starting to be far from the unit circle. When building a reduced order model with binning, only modes of relatively low frequency will be present for longer than a few timesteps. The RBFs are affected less severely, although many modes still

have a small amplitude. Finally, at the highest downsampling factor, only few modes at low frequency are close to the unit circle for both binning and the RBFs. Even though the RBFs match the POD amplitudes of the gridded case for a larger number of modes, this appears to be an insufficient condition to have high-frequency eigenvalues close to the unit circle. Spectral analysis of these correctly matched POD modes showed they indeed mostly contained small frequencies while high-frequency content is captured by neither binning nor the RBF POD and hence, also not in the DMD.

Finally, regarding the spatial structures of the DMD, Figure 11 shows the real part of spatial modes marked by the large red dots in Figure 10. These are around $St = 0.15$ for the jet and $St = 1.47$ for the airfoil; both are dominant frequencies in the spectra of the temporal correlation matrix in Figure 5. As before for the POD, the upper and lower part show the jet and




Figure 10: Eigenvalue distribution of the reduced propagator from (21) in the complex plane for the jet (left) and the airfoil (right). The top left of each side shows the gridded case with 200 and 2000 POD modes for the jet and the airfoil. The center and bottom row correspond to binning and RBFs and each column contains one downsampling factor according to the title of the subpanel. The red dots indicate the eigenvalue of a temporal structure analyzed in Fig. 11.

the airfoil, the rows the method according to the label, and the columns the downsampling factor k .

For the jet, we can make similar observations as for the spatial structures of the POD. At the lowest k , fine-scale structures are recovered well. In the developing wall jet, the vortices are clearly visible for both binning and RBF. A frequency shift is observed compared to the gridded case, but this is simply a matter of also considering the imaginary part. As k increases, these structures are diminished, and larger structures in the free jet are smeared out for binning. Since the DMD produces only one frequency per mode, there is no spectral mixing; instead, small structures near the wall are smoothed out. This effect is not as extreme for the RBFs, as fine structures near the wall are still visible at $k = 16$. The identified frequency of this mode varies little, with a maximum deviation of $\Delta St = 0.002$ for binning at $k = 16$. Other dominant structures are found equally well, with similar, small differences in frequency.

For the airfoil, the same observations hold. At low k , all three methods are almost indistinguishable, and the small structures of the separation bubble are recovered well. As k increases, both methods start to suffer, with binning suffering more than the RBFs since the spatial low-pass filtering of the method translates into smoothed spatial structures. Also, the difference in the detected frequency is small across both methods and all values of the downsampling factor k .

In summary, since the DMD is computed through the POD, it inherits some of the faults: high frequencies are diminished

in the POD spectra for larger values of k and thus, also in the DMD. The spatial structures likewise are affected by the spatial smoothing of the data mapping method, leading to smoothed structures. Spectral mixing is alleviated since the DMD gives only one frequency per mode

5. Conclusion and Perspectives

We propose a meshless framework to compute data-driven modal decompositions of scattered data. The method uses physics-constrained radial basis functions (RBFs) to obtain an analytical expression of the data at every time instance. This expression is used to compute the inner products required for every modal decomposition, namely the proper orthogonal decomposition (POD), the dynamic mode decomposition (DMD), the two spectral PODs (SPODs), and multi-scale POD (mPOD). All inner products in space of the data are computed through cheap Monte Carlo sampling of the basis functions, greatly reducing computational cost and allowing reuse of expressions between the decompositions. The approach results in super-resolution of all spatial structures, meaning they can be expressed in an arbitrary set of points.

We compare our proposed method with binning, a simple moving average over a discrete set of points that approximates the spatial inner product with summations. We evaluate both methods on two different test cases: an impinging jet on a plate from particle image velocimetry (PIV) and a flow past a transitional airfoil from large eddy simulations (LES). Both




Figure 11: Spatial structures of a dominant frequency marked with a red dot in Fig. 10 for the jet (top) and the airfoil (bottom). The top left subpanel each shows the ground truth from the gridded data and the rows correspond to the method labeled on the left. The columns correspond to increasing downsampling factors according to the title of each subpanel and the range of contours is shared across the panels of each case. In each subpanel, the frequency of each mode is also given. For the airfoil, the aspect ratio is slightly increased for better visualization.

cases feature a wide range of scales. Yet, the temporal and spatial resolution of the PIV is, typical for experiments, coarse, while both are finer for the LES. For both datasets, we generate scattered data with different spatial downsamplings to compare how binning and RBF-based inner products handle varying data densities.

For binning, the temporal correlation matrix appears smoothed and dampened, and its spectra show diminished higher frequencies. Instead, the RBFs almost perfectly match the PIV and LES in the relevant range of frequencies, even when reducing the original data density by factors of 16 or larger. These effects are much stronger for the jet, where the temporal resolution is limited; for the airfoil, the high temporal resolution is able to mitigate the loss of spatial information.

Nevertheless, the performance of the meshless

decomposition framework is inherently tied to the quality of the RBF regression and thus to the spatial sampling density. The results presented here are obtained under conditions where the seeding density is sufficiently high for the RBF approximation to resolve the relevant spatial gradients and thereby yield more accurate inner products, spectra, and modal shapes. Although the RBF approach is known to systematically outperform binning as the seeding density decreases ([58]), a reduction in seeding inevitably reduces its advantage.

These observations translate from the eigenvalue decompositions of the correlation matrix for the POD, mPOD, and SPOD [19] into the modal amplitudes and temporal structures. Also for the Fourier-based SPOD [20], the RBF-based inner products outperform binning and give more accurate spectra for both the PIV and the LES. Binning always

attenuates higher frequency content if temporal and spatial resolution are insufficient. In these conditions, we also observe increased spectral mixing in the POD modes for binning while the RBFs remain largely unaffected. An mPOD with filter banks at high frequencies or a Sieber SPOD with large filter frequencies already resolves this issue.

The spatial structures are computed from the projection of the data onto the temporal matrix. Therefore, if the temporal structures are affected by the spectral mixing, then so are the spatial structures. Also, since the data incorporates any low-pass filtering of the method, this also translates into the spatial structures. Thus, even if a perfect amplitude and temporal structure are computed, binning will always show low-pass filtering, which the RBFs do not since regression is a more sophisticated method. These observations are consistent across all of the mentioned POD variants.

The DMD also inherits these properties since it is computed from a truncated POD. If the data are downsampled, higher frequency content is attenuated in the temporal structures of the POD and also the eigenvalues of the reduced propagator, which are located further from the unit circle. This occurs more severely for high-frequency modes and more for binning than for the RBFs. The spatial structures are computed from the POD and thus suffer the same fate: If the temporal structure is poor, then so is the spatial one, and any spatial filtering in the data mapping method also translates into the spatial structure. Since the RBFs can better fill the gaps of scattered data, they naturally cope with this better.

The most natural path for future research is to also include temporal information. For binning, no off-the-shelf method exists, but 4D extensions with a Gaussian or binomial weight seem a feasible way to incorporate spatio-temporal smoothing. For the RBF regression, temporal information has been used by adding temporal super-resolution of Lagrangian tracks [51] and iterative updating between tracks and snapshots [57]. An analytical expression of the data in time would also allow to compute the inner products in time when the data are projected onto the temporal modes.

Acknowledgements

M.R. is funded by the Fonds de la Recherche Scientifique (F.R.S.-FNRS) through the FRIA grant No. FC57471. M.A.M.'s contributions are supported by the European Research Council (ERC) under the Horizon Europe programme (Starting Grant No. 101165479, RE-TWIST). The data are visualized using perceptually uniform colormaps to prevent distortion of the data and exclusion of readers with color-vision deficiencies [81, 82].

References

[1] G. Berkooz, P. Holmes, J. L. Lumley, The Proper Orthogonal Decomposition in the Analysis of Turbulent Flows, *Annual Review of Fluid Mechanics* 25 (1993) 539–575. doi:10.1146/annurev.fl.25.010193.002543.

[2] C. W. Rowley, S. T. M. Dawson, Model Reduction for Flow Analysis and Control, *Annual Review of Fluid Mechanics* 49 (2017) 387–417. doi:10.1146/annurev-fluid-010816-060042.

[3] K. Taira, S. L. Brunton, S. T. M. Dawson, C. W. Rowley, T. Colonius, B. J. McKeon, O. T. Schmidt, S. Gordeyev, V. Theofilis, L. S. Ukeiley, Modal Analysis of Fluid Flows: An Overview, *AIAA Journal* 55 (2017) 4013–4041. doi:10.2514/1.J056060.

[4] M. Mendez, M. Raiola, A. Masullo, S. Discetti, A. Ianiro, R. Theunissen, J.-M. Buchlin, POD-based Background Removal for Particle Image Velocimetry, *Experimental Thermal and Fluid Science* 80 (2017) 181–192. doi:10.1016/j.expthermflusci.2016.08.021.

[5] M. Kirby, L. Sirovich, Application of the karhunen-loeve procedure for the characterization of human faces, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12 (1990) 103–108. doi:10.1109/34.41390.

[6] A. Procacci, M. M. Kamal, M. A. Mendez, S. Hochgreb, A. Coussement, A. Parente, Multi-scale proper orthogonal decomposition analysis of instabilities in swirled and stratified flames, *Physics of Fluids* 34 (2022). doi:10.1063/5.0127956.

[7] S. D. Fassois, J. S. Sakellariou, Time-series methods for fault detection and identification in vibrating structures, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 365 (2007) 411–448. doi:10.1098/rsta.2006.1929.

[8] B. F. Feeny, R. Kappagantu, On the physical interpretation of proper orthogonal modes in vibrations, *Journal of Sound and Vibration* 211 (1998) 607–616. doi:10.1006/jjsvi.1997.1386.

[9] N. F. Troje, Decomposing biological motion: A framework for analysis and synthesis of human gait patterns, *Journal of Vision* 2 (2002) 371–387. doi:10.1167/2.5.2.

[10] M. Nonnenmacher, S. C. Turaga, J. H. Macke, Extracting low-dimensional dynamics from multiple large-scale neural population recordings by learning to predict correlations, in: 31st Conference on Neural Information Processing Systems, 2017, pp. 1–11.

[11] R. X. Shi, J. F. Wang, C. D. Xu, S. J. Lai, W. Z. Yang, Spatiotemporal pattern of hand-foot-mouth disease in China: an analysis of empirical orthogonal functions, *Public Health* 128 (2014) 367–375. doi:10.1016/j.puhe.2014.01.005.

[12] M. Raffel, C. E. Willert, F. Scarano, C. Kähler, S. T. Wereley, J. Kompenhans, *Particle Image Velocimetry - A Practical Guide*, 3 ed., Springer International Publishing AG, 2018. doi:10.1007/978-3-319-68852-7.

[13] R. J. LeVeque, *Finite Difference Methods for Ordinary and Partial Differential Equations: Steady-State and Time-Dependent Problems*, Society for Industrial and Applied Mathematics, 2007. doi:10.1137/1.9780898717839.

[14] M. A. Mendez, A. Ianiro, B. R. Noack, S. L. Brunton (Eds.), *Data-Driven Fluid Mechanics: Combining First Principles and Machine Learning*, Cambridge University Press, 2023. doi:10.1017/9781108896214.

[15] J. L. Lumley, The structure of inhomogeneous turbulent flows, *Atmospheric turbulence and radio wave propagation* (1967) 166–178.

[16] J. L. Lumley, Coherent Structures in Turbulence, in: R. E. Meyer (Ed.), *Transition and Turbulence*, Academic Press, 1981, pp. 215–242. doi:10.1016/B978-0-12-493240-1.50017-X.

[17] L. Sirovich, Turbulence and the dynamics of coherent structures. part I: Coherent structures, *Quarterly of Applied Mathematics* 45 (1987) 561–571.

[18] P. J. Schmid, Dynamic mode decomposition of numerical and experimental data, *Journal of Fluid Mechanics* 656 (2010) 5–28. doi:10.1017/S0022112010001217.

[19] M. Sieber, C. O. Paschereit, K. Oberleithner, Spectral proper orthogonal decomposition, *Journal of Fluid Mechanics* 792 (2016) 798–828. doi:10.1017/jfm.2016.103.

[20] A. Towne, O. T. Schmidt, T. Colonius, Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis, *Journal of Fluid Mechanics* 847 (2018) 821–867. doi:10.1017/jfm.2018.283.

- [21] M. A. Mendez, M. Balabane, J.-M. Buchlin, Multi-scale proper orthogonal decomposition of complex fluid flows, *Journal of Fluid Mechanics* 870 (2019) 988–1036. doi:[10.1017/jfm.2019.212](https://doi.org/10.1017/jfm.2019.212).
- [22] M. A. Mendez, D. Hess, B. B. Watz, J.-M. Buchlin, Multiscale proper orthogonal decomposition (mPOD) of TR-PTV data—a case study on stationary and transient cylinder wake flows, *Measurement Science and Technology* 31 (2020) 094014. doi:[10.1088/1361-6501/ab82be](https://doi.org/10.1088/1361-6501/ab82be).
- [23] M. L. Stein, *Interpolation of Spatial Data: Some Theory for Kriging*, Springer, New York, 1999.
- [24] G. E. Fasshauer, *Meshfree approximation methods with MATLAB*, volume 6, World Scientific, 2007.
- [25] E. J. Candès, B. Recht, Exact Matrix Completion via Convex Optimization, *Foundations of Computational Mathematics* 9 (2009) 717–772. doi:[10.1007/s10208-009-9045-5](https://doi.org/10.1007/s10208-009-9045-5).
- [26] M. A. Oliver, R. Webster, Kriging: a method of interpolation for geographical information systems, *International journal of geographical information systems* 4 (1990) 313–332. doi:[10.1080/02693799008941549](https://doi.org/10.1080/02693799008941549).
- [27] R. Everson, L. Sirovich, Karhunen–Loève procedure for gappy data, *Journal of the Optical Society of America A* 12 (1995) 1657–1664. doi:[10.1364/JOSA.A.12.001657](https://doi.org/10.1364/JOSA.A.12.001657).
- [28] H. Gunes, S. Sirisup, G. Karniadakis, Gappy data: To krig or not to krig?, *Journal of Computational Physics* 212 (2006) 358–382. doi:[10.1016/j.jcp.2005.06.023](https://doi.org/10.1016/j.jcp.2005.06.023).
- [29] A. Nekkanti, O. T. Schmidt, Gappy spectral proper orthogonal decomposition, *Journal of Computational Physics* 478 (2023) 111950–111950. doi:[10.1016/j.jcp.2023.111950](https://doi.org/10.1016/j.jcp.2023.111950).
- [30] K. Menon, R. Mittal, Dynamic mode decomposition based analysis of flow over a sinusoidally pitching airfoil, *Journal of Fluids and Structures* 94 (2020) 102886. doi:[10.1016/j.jfluidstructs.2020.102886](https://doi.org/10.1016/j.jfluidstructs.2020.102886).
- [31] B. Li, J. Garicano-Mena, E. Valero, A dynamic mode decomposition technique for the analysis of non-uniformly sampled flow data, *Journal of Computational Physics* 468 (2022) 111495. doi:[10.1016/j.jcp.2022.111495](https://doi.org/10.1016/j.jcp.2022.111495).
- [32] E. Smith, I. Variansyah, R. McClarren, Variable Dynamic Mode Decomposition for Estimating Time Eigenvalues in Nuclear Systems, *Nuclear Science and Engineering* 197 (2023) 1769–1778. doi:[10.1080/00295639.2022.2142025](https://doi.org/10.1080/00295639.2022.2142025).
- [33] D. Donoho, Compressed sensing, *IEEE Transactions on Information Theory* 52 (2006) 1289–1306. doi:[10.1109/tit.2006.871582](https://doi.org/10.1109/tit.2006.871582).
- [34] S. L. Brunton, J. L. Proctor, J. H. Tu, J. N. Kutz, Compressed sensing and dynamic mode decomposition, *Journal of Computational Dynamics* 2 (2015) 165–191. doi:[10.3934/jcd.2015002](https://doi.org/10.3934/jcd.2015002).
- [35] J. Matulis, H. Bindra, Thermal field reconstruction and compressive sensing using proper orthogonal decomposition, *Frontiers in Energy Research* 12 (2024) 1336540. doi:[10.3389/fenrg.2024.1336540](https://doi.org/10.3389/fenrg.2024.1336540).
- [36] M. Xiong, L. Chen, J. Ming, J. Shin, Accelerating the bayesian inference of inverse problems by using data-driven compressive sensing method based on proper orthogonal decomposition, *Electronic Research Archive* 29 (2021) 3383–3403. doi:[10.3934/era.2021044](https://doi.org/10.3934/era.2021044).
- [37] M. F. Fathi, A. Bakhshinejad, A. Baghaie, R. M. D’Souza, Dynamic Denoising and Gappy Data Reconstruction Based on Dynamic Mode Decomposition and Discrete Cosine Transform, *Applied Sciences* 8 (2018). doi:[10.3390/app8091515](https://doi.org/10.3390/app8091515).
- [38] J. J. Monaghan, Smoothed particle hydrodynamics, *Annual Review of Astronomy and Astrophysics* 30 (1992) 543–574.
- [39] A. Leonard, Vortex methods for flow simulation, *Journal of Computational Physics* 37 (1980) 289–335. doi:[10.1016/0021-9991\(80\)90040-6](https://doi.org/10.1016/0021-9991(80)90040-6).
- [40] N. T. Ouellette, H. Xu, E. Bodenschatz, A quantitative study of three-dimensional Lagrangian particle tracking algorithms, *Experiments in Fluids* 40 (2006) 301–313. doi:[10.1007/s00348-005-0068-7](https://doi.org/10.1007/s00348-005-0068-7).
- [41] D. Schanz, S. Gesemann, A. Schröder, Shake-The-Box: Lagrangian particle tracking at high particle image densities, *Experiments in Fluids* 57 (2016). doi:[10.1007/s00348-016-2157-1](https://doi.org/10.1007/s00348-016-2157-1).
- [42] S. Tan, A. Salibindla, A. U. M. Masuk, R. Ni, Introducing OpenLPT: new method of removing ghost particles and high-concentration particle shadow tracking, *Experiments in Fluids* 61 (2020). doi:[10.1007/s00348-019-2875-2](https://doi.org/10.1007/s00348-019-2875-2).
- [43] A. Schröder, D. Schanz, 3D Lagrangian Particle Tracking in Fluid Mechanics, *Annual Review of Fluid Mechanics* 55 (2023) 511–40. doi:[10.1146/annurev-fluid-031822-041721](https://doi.org/10.1146/annurev-fluid-031822-041721).
- [44] J. C. Agüí, J. Jiménez, On the performance of particle tracking, *Journal of Fluid Mechanics* 185 (1987) 447–468. doi:[10.1017/S0022112087003252](https://doi.org/10.1017/S0022112087003252).
- [45] J. Schobesberger, D. Worf, P. Lichtner, S. Yücesan, C. Hauer, H. Habersack, C. Sindelar, Role of low-order proper orthogonal decomposition modes and large-scale coherent structures on sediment particle entrainment, *Journal of Hydraulic Research* 60 (2021). doi:[10.1080/00221686.2020.1869604](https://doi.org/10.1080/00221686.2020.1869604).
- [46] V. Springel, Smoothed Particle Hydrodynamics in Astrophysics, *Annual Review of Astronomy and Astrophysics* 48 (2010) 391–430. doi:[10.1146/annurev-astro-081309-130914](https://doi.org/10.1146/annurev-astro-081309-130914).
- [47] X. Wu, H. Saaïd, J. Voorneveld, T. Claessens, J. Westenberg, N. Jong, J. G. Bosch, S. Kenjeres, 4d flow patterns and relative pressure distribution in a left ventricle model by shake-the-box and proper orthogonal decomposition analysis, *Cardiovascular Engineering and Technology* 14 (2023). doi:[10.1007/s13239-023-00684-0](https://doi.org/10.1007/s13239-023-00684-0).
- [48] C. Marshall, R. Dorrell, G. Keevil, J. Peakall, S. Tobias, On the role of transverse motion in pseudo-steady gravity currents, *Experiments in Fluids* 64 (2023). doi:[10.1007/s00348-023-03599-7](https://doi.org/10.1007/s00348-023-03599-7).
- [49] J. Cortina-Fernández, C. Sanmiguel Vila, A. Ianiro, S. Discetti, From sparse data to high-resolution fields: ensemble particle modes as a basis for high-resolution flow characterization, *Experimental Thermal and Fluid Science* 120 (2021) 110178. doi:[10.1016/j.expthermflusci.2020.110178](https://doi.org/10.1016/j.expthermflusci.2020.110178).
- [50] A. Grille Guerra, A. Sciacchitano, F. Scarano, Iterative modal reconstruction for sparse particle tracking data, *Physics of Fluids* 36 (2024) 075107. doi:[10.1063/5.0209527](https://doi.org/10.1063/5.0209527).
- [51] S. Gesemann, F. Huhn, D. Schanz, A. Schröder, From Noisy Particle Tracks to Velocity, Acceleration and Pressure Fields using B-splines and Penalties, in: *18th International Symposium on Applications of Laser Techniques to Fluid Mechanics*, 2016, pp. 1–17.
- [52] Y. Jeon, M. Müller, D. Michaelis, Fine scale reconstruction (VIC#) by implementing additional constraints and coarse-grid approximation into VIC+, *Experiments in Fluids* 63 (2022). doi:[10.1007/s00348-022-03422-9](https://doi.org/10.1007/s00348-022-03422-9).
- [53] A. Sciacchitano, B. Leclaire, A. Schröder, On the accuracy of data assimilation algorithms for dense flow field reconstructions, *Experiments in Fluids* 66 (2025). doi:[10.1007/s00348-025-03969-3](https://doi.org/10.1007/s00348-025-03969-3).
- [54] P. Sperotto, S. Pieraccini, M. A. Mendez, A Meshless Method to Compute Pressure Fields from Image Velocimetry, *Measurement Science and Technology* 33 (2022) 094005. doi:[10.1088/1361-6501/ac70a9](https://doi.org/10.1088/1361-6501/ac70a9).
- [55] P. Sperotto, M. Ratz, M. A. Mendez, SPICY: a Python toolbox for meshless assimilation from image velocimetry using radial basis functions, *Journal of Open Source Software* 9 (2024). doi:[10.21105/joss.05749](https://doi.org/10.21105/joss.05749).
- [56] P. Sperotto, B. Watz, D. Hess, Meshless track assimilation (MTA) of 3D PTV data, *Measurement Science and Technology* 35 (2024) 086005. doi:[10.1088/1361-6501/ad3f36](https://doi.org/10.1088/1361-6501/ad3f36).
- [57] L. Li, Z. Pan, Three-dimensional time-resolved Lagrangian flow field reconstruction based on constrained least squares and stable radial basis function, *Experiments in Fluids* 65 (2024) 57. doi:[10.1007/s00348-024-03788-y](https://doi.org/10.1007/s00348-024-03788-y).
- [58] M. Ratz, M. A. Mendez, A meshless and binless approach to compute statistics in 3D ensemble PTV, *Experiments in Fluids* 65 (2024). doi:[10.1007/s00348-024-03878-x](https://doi.org/10.1007/s00348-024-03878-x).
- [59] D. Rigutto, M. Ratz, M. A. Mendez, A meshless data-tailored approach to compute statistics from scattered data with adaptive radial basis functions, Submitted to: *Experiments in Fluids*, 2025.
- [60] X. Zhang, K. Song, M. Lu, Meshless methods based on collocation with radial basis functions, *Computational Mechanics* 26 (2000) 333–343. doi:[10.1007/s004660000181](https://doi.org/10.1007/s004660000181).

- [61] W. Chen, Z.-J. Fu, C. S. Chen, Recent Advances in Radial Basis Function Collocation Methods, 1 ed., Springer Berlin, Heidelberg, 2014. doi:[10.1007/978-3-642-39572-7](#).
- [62] B. Fornberg, N. Flyer, Solving PDEs with radial basis functions, *Acta Numerica* 24 (2015) 215–258. doi:[10.1017/S0962492914000130](#).
- [63] I. Tirelli, M. A. Mendez, A. Ianiro, S. Discetti, A meshless method to compute the proper orthogonal decomposition and its variants from scattered data, *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 481 (2025). doi:[10.1098/rspa.2024.0526](#).
- [64] J. Ramsay, B. Silverman, Principal components analysis for functional data, *Functional data analysis* (2005) 147–172. doi:[10.1007/0-387-22751-2_8](#).
- [65] J.-L. Wang, J.-M. Chiou, H.-G. Müller, A survey of functional principal component analysis, *The American Statistician* 70 (2016) 127–137. doi:[10.1080/00031305.2016.1141704](#).
- [66] P. Hall, J. L. Horowitz, On properties of functional principal components analysis, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 68 (2006) 109–126. doi:[10.1111/j.1467-9868.2005.00535.x](#).
- [67] T. Hastie, R. Tibshirani, J. Friedman, *The Elements of Statistical Learning*, Springer New York, 2009. doi:[10.1007/978-0-387-84858-7](#).
- [68] C. M. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2011.
- [69] D. C.-L. Fong, M. Saunders, LSMR: An iterative algorithm for sparse least-squares problems, *arxiv*, 2011. doi:[10.48550/arXiv.1006.0758](#).
- [70] P. Holmes, J. L. Lumley, G. Berkooz, C. W. Rowley, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, Cambridge University Press, 2012. doi:[10.1017/cbo9780511919701](#).
- [71] P. Welch, The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms, *IEEE Transactions on Audio and Electroacoustics* 15 (1967) 70–73. doi:[10.1109/TAU.1967.1161901](#).
- [72] G. H. Golub, C. F. Van Loan, *Matrix computations*, Johns Hopkins studies in the mathematical sciences, fourth edition ed., The Johns Hopkins University Press, Baltimore, 2013. *Literaturangaben und Index*.
- [73] J. H. Tu, C. W. Rowley, D. M. Luchtenburg, S. L. Brunton, J. N. Kutz, On dynamic mode decomposition: Theory and applications, *Journal of Computational Dynamics* 1 (2014) 391–421. doi:[10.3934/jcd.2014.1.391](#).
- [74] H. von Storch, J. Xu, Principal oscillation pattern analysis of the 30- to 60-day oscillation in the tropical troposphere: Part i: Definition of an index and its prediction, *Climate Dynamics* 4 (1990) 175–190. doi:[10.1007/bf00209520](#).
- [75] C. Penland, A stochastic model of Indo-Pacific sea surface temperature anomalies, *Physica D: Nonlinear Phenomena* 98 (1996) 534–558. doi:[10.1016/0167-2789\(96\)00124-8](#).
- [76] K. Hasselmann, Pips and pops: The reduction of complex dynamical systems using principal interaction and oscillation patterns, *Journal of Geophysical Research: Atmospheres* 93 (1988) 11015–11021. doi:[10.1029/jd093id09p11015](#).
- [77] M. Mendez, R. Poletti, L. Schena, *Fundamentals of Dimensionality Reduction*, Cambridge University Press, 2025, pp. 259–299.
- [78] I. Tirelli, A. Ianiro, S. Discetti, A simple trick to improve the accuracy of PIV/PTV data, *Experimental Thermal and Fluid Science* 145 (2023) 110872. doi:[10.1016/j.expthermflusci.2023.110872](#).
- [79] C.-A. Yeh, K. Taira, Resolvent-analysis-based design of airfoil separation control, *Journal of Fluid Mechanics* 867 (2019) 572–610. doi:[10.1017/jfm.2019.163](#).
- [80] A. Towne, S. T. M. Dawson, G. A. Brès, A. Lozano-Durán, T. Saxton-Fox, A. Parthasarathy, A. R. Jones, H. Biler, C.-A. Yeh, H. D. Patel, K. Taira, A Database for Reduced-Complexity Modeling of Fluid Flows, *AIAA Journal* 61 (2023) 2867–2892. doi:[10.2514/1.J062203](#).
- [81] F. Crameri, G. E. Shephard, P. J. Heron, The misuse of colour in science communication, *Nature Communications* 11 (2020). doi:[10.1038/s41467-020-19160-7](#).
- [82] F. Crameri, *Scientific colour maps*, 2023. doi:[10.5281/ZENODO.8409685](#).

Appendix A. Nomenclature

We use lowercase letters for scalar continuous quantities, e.g. $t \in \mathbb{R}$. Bold lowercase letters are used for vector valued continuous quantities, e.g. $\mathbf{x} \in \mathbb{R}^d$. We use the roman bold symbol to denote vectors collecting samples of a scalar quantity, e.g. $\mathbf{t}$ or $\mathbf{c}$, and the subscript $\mathbf{t}_k$ or $\mathbf{c}_m$ denote the k -th or the m -th entry of these vectors. Similarly, a collection of vector valued quantities is indicated with a roman bold capital symbol, hence $\mathbf{X} \in \mathbb{R}^{d \times n}$ is a collection of n vectors of size d . The subscript on a matrix is generally used to denote its column, hence $\mathbf{X}_m \in \mathbb{R}^d$ is the m -th column of $\mathbf{X}$, thus the m -th vector valued sample. However, sometimes a subscript is used to distinguish variables: for example $\mathbf{X}_p$ is the set of training points while $\mathbf{X}_c$ is the set of collocation points. In this case, the sampling is always used as a second script, e.g. $\mathbf{X}_{c,m}$ is the m -th collocation point. Confusion in the usage of subscripts for different purposes can be easily avoided by noticing that only seven letters are used for indices: these are i, j, r, s, l, m, n . Any other subscript (such as p or c) is thus not an index.

We use the square brackets to denote vectors or matrices created from collections of scalars or vectors. Therefore, $\mathbf{t} = [\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_{n_t}] = \{\mathbf{t}_i\}_{i=1}^{n_t}$ is the vector collecting the times of all spaced samples in time while $\mathbf{X} = [\mathbf{X}_1, \mathbf{X}_2 \dots \mathbf{X}_{n_p}] = \{\mathbf{X}_n\}_{n=1}^{n_p}$ is a matrix collecting n_p location in space $\mathbf{X}_n$. In either case, we use the curly brackets $\{\}$ for sequences of scalars or vectors.

Equivalently, when stress is placed on the specific scalars forming a vector or vectors forming a matrix, a superscript with parenthesis is used. For example $\hat{\mathbf{W}}_{w,u} = [\hat{\mathbf{w}}_u^{(1)}, \hat{\mathbf{w}}_u^{(2)} \dots \hat{\mathbf{w}}_u^{(n_B)}]$ is the matrix built by stacking the DFT of the weights $\mathbf{w}_u$, denoted as $\hat{\mathbf{w}}_u$, at each of the n_B blocks of the Welch averaging procedure.

Concerning the indices of entries in vectors and matrices, we adopt a Python-like notation. Specifically, $A[i, j]$ denotes the entry in the i -th row and j -th column of the matrix A , while $\mathbf{v}[i]$ refers to the i -th entry of a vector $\mathbf{v}$. Accordingly, $A[i, j] = A_j[i]$ are two equivalent ways of referencing the same entry. We retain both notations: the former emphasizes how certain matrices are constructed, while the latter highlights the role of certain column vectors. Moreover, since vectors and matrix columns typically collect samples in time, space, or frequency domains, index and function notations are both relevant and used. Thus the notation $\hat{u}(\mathbf{x}, \mathbf{f}_l)$ emphasizes that $\hat{u}$ is a continuous function in $\mathbf{x}$ but a discrete function in $\mathbf{f}_l$ and the column $\mathbf{W}_{u,m} = \mathbf{w}_u(\mathbf{t}_m)$ collects the RBF weight vector at the m -th time step.