CrossMark

RESEARCH ARTICLE

Solenoidal filtering of volumetric velocity measurements using Gaussian process regression

Iliass Azijli[1] · Richard P. Dwight[1]

Abstract Volumetric velocity measurements of incompressible flows contain spurious divergence due to measurement noise, despite mass conservation dictating that the velocity field must be divergence-free (solenoidal). We investigate the use of Gaussian process regression to filter spurious divergence, returning analytically solenoidal velocity fields. We denote the filter solenoidal Gaussian process regression (SGPR) and formulate it within the Bayesian framework to allow a natural inclusion of measurement uncertainty. To enable efficient handling of large data sets on regular and near-regular grids, we propose a solution procedure that exploits the Toeplitz structure of the system matrix. We apply SGPR to two synthetic and two experimental test cases and compare it with two other recently proposed solenoidal filters. For the synthetic test cases, we find that SGPR consistently returns more accurate velocity, vorticity and pressure fields. From the experimental test cases, we draw two important conclusions. Firstly, it is found that including an accurate model for the local measurement uncertainty further improves the accuracy of the velocity field reconstructed with SGPR. Secondly, it is found that all solenoidal filters result in an improved reconstruction of the pressure field, as verified with microphone measurements. The results obtained with SGPR are insensitive to correlation length, demonstrating the robustness of the filter to its parameters.

✉ Iliass Azijli
i.azijli@tudelft.nl

Richard P. Dwight
r.p.dwight@tudelft.nl

[1] TU Delft, Kluyverweg 1, 2629 HS Delft, The Netherlands

1 Introduction

A number of experimental velocimetry techniques have recently become available that allow the determination of the three velocity components (3C) in a volume (3D) of fluid. In particular, there are 3D particle tracking velocimetry (PTV) (Maas et al. 1993), scanning particle image velocimetry (PIV) (Brücker 1995), holographic PIV (Hinsch 2002), magnetic resonance velocimetry (MRV) (Elkins et al. 2003) and most recently tomographic PIV (Elsinga et al. 2006). Since the mentioned techniques measure the three velocity components in a volume, they enable one to calculate the full velocity gradient tensor. Not only does this allow a less ambiguous analysis of coherent structures than in the case where only 2D velocity slices are available, but it potentially allows accurate inference of pressure fields via the Navier-Stokes equations (van Oudheusden 2013). In the future, this latter capability is expected to have important applications in non-intrusive evaluation of forces, and the study of acoustic sources (Scarano 2013).

For incompressible flows, the trace of the velocity gradient tensor is zero by the mass conservation equation, i.e., the velocity field is solenoidal (divergence-free). However, measurement errors will cause spurious nonzero divergence in measured fields. It is possible to reduce this spurious divergence somewhat through ad hoc smoothing the data, as shown by Zhang et al. (1997). However, this results in a loss of resolution, without achieving a perfectly divergence-free velocity field. A more physically consistent approach is to enforce mass conservation in the filter itself.

Including this constraint can be particularly advantageous when calculating physics-based derived quantities. For example, Sadati et al. (2011) found a significant improvement in the calculation of stress and optical signal

fields from velocity measurements of a viscoelastic flow. In this paper, we demonstrate improvements in the calculation of pressure fields. Additionally, it is essential when the measurements are to be used as input for numerical simulations, e.g., for temporal flow reconstruction. In particular, in the field of PIV, Sciacchitano et al. (2012) proposed a method that uses the unsteady incompressible Navier–Stokes equations to fill spatial gaps. Similarly Schneiders et al. (2014) proposed a vortex-in-cell method to fill temporal gaps. These methods currently build a solenoidal velocity field from measurement data in an ad hoc way, without accounting for measurement noise or flow scales. This introduces an error in their input which pollutes the temporal reconstruction. Both methods will benefit from a preprocessing step that handles the noisy data and imposes mass conservation consistently.

1.1 Solenoidal filters

A number of methods have recently been proposed in the literature to filter out spurious divergence from 3C-3D velocity measurements. We have identified three general approaches, based on different principles.

Helmholtz representation theorem The first approach is based on the Helmholtz representation theorem. This states that any finite, twice differentiable vector field **u** can be decomposed into a solenoidal vector field $\mathbf{u}_{sol}$ plus an irrotational vector field $\mathbf{u}_{irrot}$ (Segel 2007):

$$\mathbf{u} = \mathbf{u}_{sol} + \mathbf{u}_{irrot} = \nabla \times \mathbf{a} + \nabla \psi, \qquad (1)$$

where **a** is a vector potential and ψ is a scalar potential. Taking the divergence on both sides of Eq. 1 and applying $\nabla \cdot \mathbf{u}_{sol} = 0$ gives a Poisson equation:

$$\nabla^2 \psi = \nabla \cdot \mathbf{u}. \qquad (2)$$

Solving Eq. 2 gives ψ, from which the solenoidal velocity field can be obtained using Eq. 1. Difficulties arise in the specification of the boundary conditions, as in the absence of walls the boundary conditions are unknown. One extreme is to specify $\nabla \psi \cdot \mathbf{n} = \mathbf{u} \cdot \mathbf{n}$, i.e., the irrotational velocity normal to the boundary is equal to the measured velocity. As a consequence, the solenoidal velocity field normal to the boundary will be zero. The other extreme is to specify $\nabla \psi \cdot \mathbf{n} = 0$; the irrotational velocity normal to the boundary is zero, meaning that the solenoidal velocity normal to the boundary is equal to the measured velocity. This will result in an ill-posed PDE if the measured velocity field is not globally mass conserving. An additional complication arises when the vorticity field is to be calculated, because the vorticity of the solenoidal velocity is equal to that of the measured velocity (an irrotational velocity field has vanishing curl). Since the measured velocity field is

contaminated by noise, the vorticity field of the solenoidal field will not be improved. A number of authors have followed the approach of the Helmholtz representation theorem in filtering velocity measurements (Song et al. 1993; Suzuki et al. 2009; Worth 2012; Yang et al. 1993).

Least-squares variational filters The second approach is to define a minimization problem. Liburdy and Young (1992) and Sadati et al. (2011) defined an unconstrained minimization problem, where the objective function consists of three terms: (1) the discrepancy between the measured and filtered velocity field, expressed as a sum of squared differences; (2) the smoothness of the velocity gradients; (3) the divergence of the velocity field . The latter two terms are multiplied with weighting parameters that must be supplied by the user. Sadati et al. (2011) proposed using generalized cross-validation to find optimum weights.

de Silva et al. (2013) defined a constrained minimization problem, thereby avoiding the use of weights in the objective function. They called their method the divergence correction scheme (DCS). The objective function is only the discrepancy between the measured and filtered velocity fields, expressed as a sum of squared differences. The divergence-free condition is introduced as a linear equality constraint. The method was applied to a tomographic PIV experiment, and an improvement in the prediction of turbulence statistics was shown. However, it turns out that DCS is equivalent to a discretized Helmholtz Theorem in that the vorticity of the solenoidal velocity field is equal to that of the measured velocity field. A proof of this is given in the Appendix 1. Despite this, the advantage of DCS is that it does not require the user to specify boundary conditions since no Poisson equation needs to be solved.

Reconstruction with a solenoidal basis A third approach is to use a solenoidal basis, which by construction will return solenoidal velocity fields. There are several possible choices of basis. Solenoidal wavelets were proposed by Battle and Federbush (1993). Narcowich and Ward (1994) devised solenoidal radial basis functions (RBF). Lowitzsch (2005) subsequently derived a density theorem stating that any sufficiently smooth solenoidal function can be approximated arbitrarily closely by a linear combination of solenoidal RBFs. Therefore any incompressible flow may be represented by this construction. Solenoidal wavelets and RBFs have been applied in computational fluid dynamics to solve incompressible flows (Deriaz and Perrier 2006; Schräder and Wendland 2011; Urban 1996; Wendland 2009). Applications of these methods to reconstruct solenoidal velocity fields from velocity measurements have been reported as well (Busch et al. 2013; Ko et al. 2000; Vennell and Beatson 2009). A method closely related to RBF interpolation is *Kriging* or *Gaussian process regression* (GPR) (Rasmussen and Williams 2006). In this context, the radial basis

functions are referred to as covariance functions because they represent the covariance between spatial locations, since the underlying field is assumed to come from a stochastic process. Scheuerer and Schlather (2012) derived a solenoidal covariance function for Kriging. In the following, the solenoidal version of Kriging will be referred to as SGPR (solenoidal Gaussian process regression).

Another choice of basis is described in the work by Schiavazzi et al. (2014), who introduced a method, referred to in the following as solenoidal waveform reconstruction (SWR), where vortices around the edges of the measurement grid are placed, ensuring a solenoidal velocity field on a finite-volume level. It works as follows: After creating a Voronoi tessellation associated with the measurement grid, the measured velocities are converted to face fluxes. Vortices are defined around the face edges, resulting in an ill-posed linear system for their strengths. This system is solved in a least-squares sense. Once the vortex strengths have been found, the new face fluxes are converted back to nodal velocities. The basis functions used in this method are the smallest possible solenoidal waveforms compatible with the measurement grid. The advantage of this is that large local measurement errors only disturb reconstructed velocities at neighboring grid points. The disadvantage of this minimal support is that it becomes difficult to reconstruct gappy fields, i.e., regions with missing data.

1.2 Contributions

A first contribution of this paper is to derive SGPR from a Bayesian perspective, which is the subject of Sect. 2. We use this perspective since it provides a natural framework to fuse prior (physical) knowledge, in this case mass conservation, with observations while taking into account the uncertainties that arise from both worlds. Using this framework is particularly apposite considering recently developed tools that allow the a posteriori quantification of measurement uncertainty from PIV (Charonko and Vlachos 2013; Sciacchitano et al. 2013; Timmins et al. 2012; Wieneke and Prevost 2014).

A second contribution is implementing SGPR efficiently to handle large data sets, which is the subject of Sect. 3. Gaussian process regression is known to be an expensive method. This is particularly important for the present application where the data sets are large. In addition, as will be shown in the next section, including the divergence-free condition results in a larger system of equations that needs to be solved. For data sets on regular and near-regular grids as is the case for PIV, the system matrix is seen to have a multilevel block-Toeplitz structure. The Toeplitz structure allows speeding up matrix-vector multiplications, required by the conjugate gradient method, from $O(N^2)$ computations

to $O(N \log N)$ computations through the use of fast Fourier transforms (FFTs). The block structure allows parallelization of these multiplications. By using an iterative solver instead of a direct solver, these speedups also hold when spatially varying measurement uncertainty is included.

The third contribution is comparing SGPR with two other recently proposed solenoidal filters, namely DCS and SWR. From this follow the final two contributions: showing how applying a solenoidal filter results in more accurate PIV-based pressure reconstructions; showing how the inclusion of local measurement uncertainty with SGPR results in a divergence-free velocity field that follows the measurements more faithfully. DSC and SWR are described in more detail in the Appendices 1 and 2, in addition to their specific implementation in the present paper. Section 4 applies the filters to synthetic test cases. Real data from two tomographic PIV experiments are used in Sect. 5. Conclusions are drawn in Sect. 6.

2 Solenoidal Gaussian process regression

This section starts with an introduction to Gaussian process regression (GPR) using a Bayesian perspective, followed by a simple one-dimensional example to illustrate the importance of including accurate, local measurement uncertainty. We then describe how mass conservation can be enforced in the prior, resulting in solenoidal Gaussian process regression (SGPR) and we finally show that it is a generalization of solenoidal radial basis function interpolation.

2.1 Gaussian process regression: a Bayesian perspective

Consider an unobservable Gaussian process with mean μ and covariance function ϕ (Rasmussen and Williams 2006):

$$\mathcal{F} \sim \mathcal{GP}\Big(\mu(\mathbf{x}), \phi\Big(\mathbf{x}, \mathbf{x}^i; \boldsymbol{\chi}\Big)\Big),$$

where $\mathbf{x} \in \mathbb{R}^3$ is the spatial coordinate and $\boldsymbol{\chi} \in \mathbb{R}^q$ are hyperparameters. Discretizing the process at n spatial locations gives a multivariate normal random-variable

$$\mathbf{F} \sim \mathcal{N}(\boldsymbol{\mu}, P),$$
$$\mu_i = \mu\Big(\mathbf{x}^i\Big),$$
$$P^{ij} = \phi\Big(\mathbf{x}^i, \mathbf{x}^j; \boldsymbol{\chi}\Big), \quad i,j = 1, \ldots, n.$$

Its realization is the state vector $\mathbf{f} \in \mathbb{R}^n$. We assume a stationary process, i.e., the correlation between two points $\mathbf{x}^i$ and $\mathbf{x}^j$ depends only on their separation, $\mathbf{x}^i - \mathbf{x}^j$, and not their absolute positions. The covariance function then

simplifies to $\phi(r^{ij})$, where r^{ij} is the Euclidian distance between the points $\mathbf{x}^i$ and $\mathbf{x}^j$, defined as:

$$(r^{ij})^2 = \sum_{k=1}^{3} \left(\frac{x_k^i - x_k^j}{\gamma_k} \right)^2, \tag{3}$$

where parameter γ_k is the correlation length in the direction k. This will be the prior.

The statistical model for the observations is defined as $\mathbf{y} = H\mathbf{f} + \boldsymbol{\epsilon}$, where $\mathbf{y} \in \mathbb{R}^m$ contains the m observations, H is the observation matrix and $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, R)$ (assuming unbiased noise). Therefore, $\mathbf{Y}|\mathbf{f} \sim \mathcal{N}(H\mathbf{f}, R)$, the likelihood. The matrix R is known as the observation error covariance matrix, representing the measurement uncertainty. We are interested in the distribution of the true state given the observed data $p(\mathbf{f}|\mathbf{y})$, i.e., the posterior. According to Bayes' rule, the posterior distribution is proportional to the prior $p(\mathbf{f})$ times the likelihood $p(\mathbf{y}|\mathbf{f})$. The posterior is therefore normally distributed as well with mean and covariance (Wikle and Berliner 2007):

$$\mathbb{E}(\mathbf{F}|\mathbf{y}) = \boldsymbol{\mu} + PH'(R + HPH')^{-1}(\mathbf{y} - H\boldsymbol{\mu}), \tag{4}$$

$$\Sigma(\mathbf{F}|\mathbf{y}) = \left(I - PH'(R + HPH')^{-1}H\right)P. \tag{5}$$

In the following, we will assume the prior mean to be a constant field and in particular we will set $\boldsymbol{\mu} = 0$. This is acceptable since the complexity of the field that can be reconstructed is contained in the prior covariance P used (Jones et al. 1998). Also, we are primarily interested in the posterior mean. Azijli et al. (2015) use the posterior covariance to propagate the uncertainty from the measured velocity field through the Navier–Stokes equations, allowing a posteriori uncertainty quantification of PIV-derived pressure fields. The term $R + HPH'$ will be referred to as the gain matrix A. In the present paper, Eq. 4 therefore reduces to:

$$\mathbb{E}(\mathbf{F}|\mathbf{y}) = PH'A^{-1}\mathbf{y}. \tag{6}$$

2.2 1D example of GPR with local error information

To illustrate the importance of including accurate, local measurement uncertainty, we present an artificial one-dimensional reconstruction problem. Consider the true function

$$h(x) := \exp\left[-54(x - 0.5)^6\right],$$

observed at eleven equally spaced measurement locations on the interval $[0, 1]$ with added noise. The measurements do not all have the same noise level (a common situation in e.g., PIV). For simplicity, we set the (co)variance of the measurement uncertainty R to a diagonal matrix (i.e., uncorrelated), where the elements on the diagonal are represented by the vector

$$\mathbf{r} = (10^{-4}, 10^{-2}, 10^{-2}, 10^{-2}, 0.25, 0.25, 0.25, \ldots$$
$$10^{-2}, 10^{-2}, 10^{-2}, 10^{-4}).$$

Notice that the measurements at the boundaries of the interval are more accurate than measurements in the center. We use a Gaussian covariance function

$$\phi_G\left(\mathbf{x}^i, \mathbf{x}^j; \gamma\right) := \exp\left[-\left(\frac{x^i - x^j}{\gamma}\right)^2\right],$$

with $\gamma = 0.15$. Equations 6 and 5 are then used to reconstruct the true function given the eleven measurements, and the (known) error covariance. Figure 1 shows the result. We have taken $n = 101$, uniformly distributed over the domain, so the function is evaluated at significantly more points than the measurement grid. H is a sparse matrix with a few 1's picking the 11 measurement points from the 101 points. We have considered three cases in terms of our prior
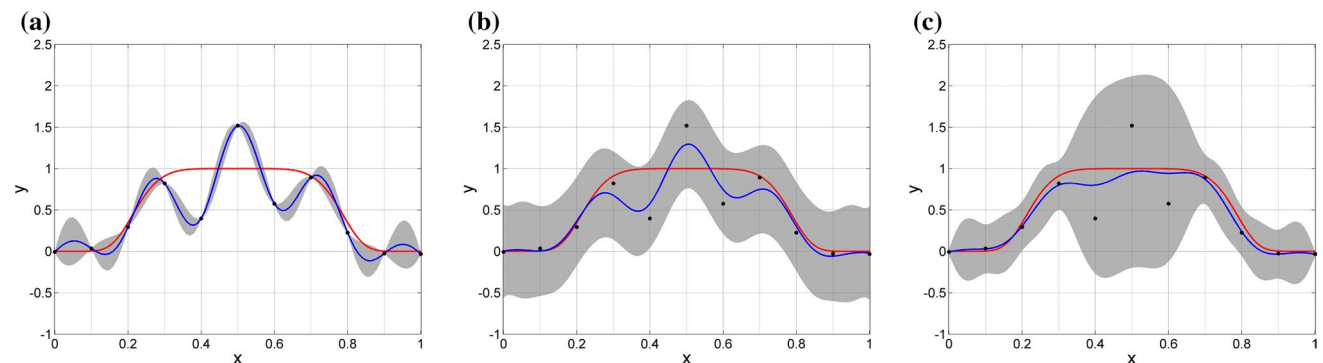

Fig. 1 Reconstructions of the function $h(x)$ (*red line*) based on eleven measurements (*black dots*), assuming perfect observations (**a**), uniform measurement uncertainties (**b**) and spatially varying meas- urement uncertainties (**c**). A reconstruction is given by a mean field (*blue lines*) with confidence intervals of ± 3 standard deviations (*gray regions*)

knowledge of the measurement uncertainty: (1) $R = 0$ (see Fig. 1a), wrongly assuming that the measurements are perfect. As a result, the reconstructed function interpolates exactly the observed points. At these locations, the posterior uncertainty is zero. Notice that due to our overconfidence in the accuracy of the measurements, the true function is well outside the ± 3 standard deviations region for a large part of the interval; (2) (see Fig. 1b), taking into account measurement uncertainty but wrongly assuming that all measurements have equal uncertainty. The reconstructed function is more accurate than in the first case. Furthermore, the posterior uncertainty does not become zero at the measurement locations. Since we assume homogeneous measurement uncertainty and the observations are equally spaced, the confidence interval is very homogeneous throughout the interval; (3) (see Fig. 1c), taking into account the correct measurement uncertainty. This gives the most accurate reconstructed function of all three cases. Notice also that the process realizes the large uncertainty in the center of the interval and as a result relies more on the prior. This is reflected by the large bounds of the confidence interval in this region.

2.3 Enforcing mass conservation

In the context of reconstructing measured velocity fields using GPR, previous works took the route of reconstructing the velocity components independently from each other, precluding enforcement of the mass conservation equation or other physical constraints (de Baar et al. 2014; Gunes and Rist 2008; Inggs and Lord 1996; Lee et al. 2008). In these works, the state is chosen as

$$\mathbf{f} = \begin{pmatrix} \mathbf{u}_1 & \mathbf{u}_2 & \mathbf{u}_3 \end{pmatrix}' \in \mathbb{R}^{3n},$$

where $\mathbf{u}_k$ are the Cartesian velocity components. The m observations are defined by $\mathbf{y} = \begin{pmatrix} \mathbf{u}_{\text{PIV},1} & \mathbf{u}_{\text{PIV},2} & \mathbf{u}_{\text{PIV},3} \end{pmatrix}' \in \mathbb{R}^{3m}$. The covariance matrix $P \in \mathbb{R}^{3n \times 3n}$ is a 3×3 block diagonal matrix, with elements P_k, $k = 1, 2, 3$ on the diagonal. The off-diagonal blocks are zero by assumption of uncorrelated velocity components.

In our case, for an incompressible flow the density is constant in space and time, giving a mass conservation equation: $\nabla \cdot \mathbf{u} = 0$. From vector calculus identities, it is known that the curl of a vector potential $\mathbf{a}$ is divergence-free. We therefore choose the state vector

$$\mathbf{f} = \begin{pmatrix} \mathbf{a}_1 & \mathbf{a}_2 & \mathbf{a}_3 \end{pmatrix}' \in \mathbb{R}^{3n},$$

where $\mathbf{a}_k$ are the Cartesian vector potential components. The observation matrix $H \in \mathbb{R}^{3m \times 3n}$, when multiplied with

the vector potential, should return its curl, therefore ensuring that the resulting vector will be divergence-free:

$$H = \begin{bmatrix} 0 & -\partial_3 \cdot & \partial_2 \cdot \\ \partial_3 \cdot & 0 & -\partial_1 \cdot \\ -\partial_2 \cdot & \partial_1 \cdot & 0 \end{bmatrix}, \tag{7}$$

where ∂_k is the partial derivative with respect to x_k. The covariance matrix $P \in \mathbb{R}^{3n \times 3n}$ is again a 3×3 block diagonal matrix, with elements P_k, $k = 1, 2, 3$ on the diagonal. However, it now represents the covariance between the components of the vector potential, not the velocity. Again, the off-diagonals of the covariance matrix are zero. But the reason for this now is that the different components of the vector potential, not the velocity, are assumed to be uncorrelated. In the most general case, there is no a priori physical knowledge to assume that there *is* a relation between them. The component HPH' of the gain matrix is equal to

$$\begin{bmatrix} \partial_{3,3}P_2 + \partial_{2,2}P_3 & -\partial_{1,2}P_3 & -\partial_{1,3}P_2 \\ -\partial_{1,2}P_3 & \partial_{3,3}P_1 + \partial_{1,1}P_3 & -\partial_{2,3}P_1 \\ -\partial_{1,3}P_2 & -\partial_{2,3}P_1 & \partial_{2,2}P_1 + \partial_{1,1}P_2 \end{bmatrix}. \tag{8}$$

It can easily be verified that the columns and rows of Eq. 8 are divergence-free. By taking the same covariance function for all directions of the vector potential ($P_1 = P_2 = P_3$) and assuming perfect measurements ($R = 0$), Eq. 8 reduces to:

$$\left(\Delta I - \nabla \nabla' \right) P_1. \tag{9}$$

This turns out to be proportional to the operator constructed by Narcowich and Ward (1994) in the context of RBF interpolation. Our method can therefore be seen as a generalization of their method that allows different covariance functions to be used for the different vector potential components, and that includes measurement uncertainty in the reconstruction via R. This latter property is a natural consequence of the Bayesian approach we followed.

Now that we have shown how SGPR is a generalization of divergence-free RBF interpolation, we exploit the latter to express how the velocity vector can be evaluated *analytically* at *any* point $\mathbf{x}$ in the field using the following expression:

$$\begin{pmatrix} u_1(\mathbf{x}) \\ u_2(\mathbf{x}) \\ u_3(\mathbf{x}) \end{pmatrix} = \sum_{i=1}^{m} \begin{bmatrix} \phi_{yy} + \phi_{zz} & -\phi_{xy} & -\phi_{xz} \\ -\phi_{xy} & \phi_{xx} + \phi_{zz} & -\phi_{yz} \\ -\phi_{xz} & -\phi_{yz} & \phi_{xx} + \phi_{yy} \end{bmatrix}_i \mathbf{c}_i, \tag{10}$$

where $\mathbf{c} = A^{-1}\mathbf{y}$ (see Eq. 6). $\phi_{kl,i}$ indicates the second partial derivative of ϕ with respect to the coordinate directions k and l, and represents the covariance between point $\mathbf{x}$ and a measurement point $\mathbf{x}^i$.

3 Efficient implementation

The practical implementation of SGPR is discussed in this section, in which the primary concern is the storage, conditioning and inversion of the (in general) dense gain matrix A. This is critical given that tomographic data sets containing 10^5–10^6 vectors are typical (Scarano 2013). By exploiting the Toeplitz structure of the gain matrix, we make the method highly efficient.

3.1 Computational cost

To evaluate the divergence-free velocity field at any point (see Eq. 10), one first has to solve the linear system $A\mathbf{c} = \mathbf{y}$, where $A \in \mathbb{R}^{3m \times 3m}$. This is the most expensive component of SGPR. The gain matrix is per definition symmetric. For scattered observations, this requires storing $3m(3m + 1)/2$ entries. If the covariance function used is positive definite, the matrix will also be positive definite (Narcowich and Ward 1994). The Cholesky factorization, with a computational complexity of $O(27m^3)$, can therefore be used as a direct solver. Alternatively, the conjugate gradient (CG) method, with a computational complexity of $O\left(9k \cdot m^2\right)$ (k is the number of iterations required to reach a given tolerance), can be used as an iterative solver. The dominant operation per iteration is the matrix-vector multiplication $A\mathbf{c}$.

For large data sets, the cost of SGPR therefore becomes prohibitive. One approach to deal with this is to use a covariance function with local support: $\phi(r \geq r_d) = 0$, where r_d is the support radius. If the support radius is much smaller than the size of the measurement domain, the gain matrix will be sparse, resulting in storage and computational savings. However, the optimum support radius may be large for a given flow field (Schaback 1995), and we do not wish to be limited by computational considerations.

If the measurements are (1) available on a regular grid, as is the case in scanning PIV and tomographic PIV; and (2) all observations have equal measurement uncertainty, then the gain matrix will have a small displacement rank (Kailath et al. 1979), allowing memory and time savings. To be more specific, the gain matrix will have a Toeplitz structure. For the present problem, it is a 3×3 block matrix where each block is a 3-level symmetric Toeplitz matrix. It is a 3-level Toeplitz matrix as we are dealing with a three-dimensional space, and the 3×3 block structure is due to the structure of Eq. 8. A 3-level symmetric Toeplitz matrix T is defined as follows:

$$T = \begin{bmatrix} T_1 & T_2 & \cdots & \cdots & T_{m_3} \\ T_2 & \ddots & \ddots & & \vdots \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \vdots & & \ddots & \ddots & T_2 \\ T_{m_3} & \cdots & \cdots & T_2 & T_1 \end{bmatrix}, \quad T_i = \begin{bmatrix} t_1 & t_2 & \cdots & \cdots & t_{m_2} \\ t_2 & \ddots & \ddots & & \vdots \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \vdots & & \ddots & \ddots & t_2 \\ t_{m_2} & \cdots & \cdots & t_2 & t_1 \end{bmatrix},$$

$$t_i = \begin{bmatrix} \tau_1 & \tau_2 & \cdots & \cdots & \tau_{m_1} \\ \tau_2 & \ddots & \ddots & & \vdots \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \vdots & & \ddots & \ddots & \tau_2 \\ \tau_{m_1} & \cdots & \cdots & \tau_2 & \tau_1 \end{bmatrix}, \tag{11}$$

where $m = m_1 m_2 m_3$ and m_k, $k = 1, 2, 3$, is the number of observations in the k-direction. The gain matrix can now be stored using 3m entries. Fast and superfast direct solvers of complexity $O\left(9m^2\right)$ and $O\left(3m \log^2 (3m)\right)$, respectively, are available (Ammar and Gragg 1988).

If the measurement uncertainty is not the same at all points, the gain matrix will have arbitrary displacement rank and no fast direct solvers are available for such a system (Ng and Pan 2010). However, when the CG method is used, accelerations can be achieved. First, the gain matrix A is decomposed into a diagonal matrix D, containing the unequal measurement uncertainties, and a 3×3 block matrix G, where each block G_i is a 3-level symmetric Toeplitz matrix:

$$A = D + \begin{bmatrix} G_1 & G_4 & G_6 \\ G_4 & G_2 & G_5 \\ G_6 & G_5 & G_3 \end{bmatrix}. \tag{12}$$

The multiplication $D\mathbf{c}$ can be carried out in $O(3m)$ operations. The matrix-vector multiplication $G\mathbf{c}$ can be accelerated by embedding each 3-level Toeplitz matrix $G_i \in \mathbb{R}^{m \times m}$ into a 3-level circulant matrix $C_i \in \mathbb{R}^{8m \times 8m}$ and then using 3D FFTs to carry out matrix-vector multiplications at a cost of $O(8m \log (8m))$ (Chan and Jin 2007). The computational complexity of the CG algorithm now becomes $O(72k \cdot m \log (8m))$. Normally, the CG method is used in combination with a preconditioner to speed up convergence. Unfortunately, Capizzano (2002) has shown that preconditioners for multilevel Toeplitz matrices are not superlinear. So in the present paper, we have implemented SGPR without a preconditioner.

Note that the assumption of observations on a regular grid should not be seen as a significant limitation. Consider two common situations where we have (1) a regular

measurement grid with missing data (e.g., from removed outliers); (2) an irregularly scattered data set that can be approximated as a subset of a regular grid. In these cases, we can set the values of the unobserved points equal to the mean of the observations (either global or local) and specify a large measurement uncertainty at these points. The size of the gain matrix increases, but this may be out weighted by the benefits of having a Toeplitz structure.

If it is not possible or efficient to exploit the Toeplitz structure, one can use the Fast Multipole Method (FMM) to speed up matrix-vector multiplications (Greengard and Rokhlin 1987), which is valid for arbitrary grids. The idea is to make use of analytical manipulations of series to achieve fast summation. The cost of a matrix-vector multiplication is similar to the FFT method, i.e., $O(N \log N)$ for matrices of size $A \in \mathbb{R}^{N \times N}$. The drawbacks of this method are that the implementation is complex and the constant in the order of its complexity may be large and thus several tricks must be used to reduce it (Darve 2000). Two notable examples of using FMM for GPR are by Billings et al. (2002) and Memarsadeghi et al. (2008).

3.2 Conditioning

The most important factors that influence the condition number of the gain matrix, $\kappa(A)$, are the covariance function, the separation distance of the data, the correlation length and the observation error (Davis and Morris 1997; Narcowich and Ward 1994). Decreasing the separation distance, increasing the correlation length and decreasing the measurement uncertainty all increase the condition number. An ill-conditioned gain matrix causes inaccurate results (in the form of numerical noise in the reconstruction) and reduces the convergence speed of the CG method. In fact, the gain matrix can become so ill-conditioned that it stops being *numerically* positive definite (i.e., in finite-precision arithmetic), even though it will always be analytically positive given a positive definite covariance function. Since our solution methods, Cholesky and CG, both rely on positive-definiteness, they cannot be applied. We therefore aim to avoid ill-conditioning by a suitable choice of the covariance function ϕ. Table 1 summarizes at which correlation length $\kappa(A)$ becomes 10^{15}. The functions $\phi_{d,k}$, $k = 1, 2, 3, 4$ are the Wendland functions with smoothness C^{2k} and spatial

Table 1 Relative correlation length γ/L at which $\kappa(A) = 10^{15}$

	$\phi_{3,1}$	$\phi_{3,2}$	$\phi_{3,3}$	$\phi_{3,4}$	ϕ_G
GPR	508.35	21.617	6.1209	3.2370	0.5175
SGPR	4.4558e9	408.63	18.503	5.4604	0.4884

The spatial dimension is 3, and the data points are defined on a regular $15 \times 15 \times 15$ grid with size $L \times L \times L$

dimension d (Wendland 2005). The Gaussian (ϕ_G), an infinitely smooth function, is defined as $\phi_G(r) = \exp\left(-\alpha^2 r^2\right)$. The constant α was set to 3.3 to make it resemble $\phi_{d,3}$ as closely as possible. Unlike the Gaussian, the Wendland functions have compact support: $\phi_{d,k} = 0$ for $r \geq 1$. The Gaussian covariance function appears to be the most ill-conditioned, a property attributed to the fact that it is infinitely differentiable (Davis and Morris 1997). The smoother the Wendland function, the more it approaches the Gaussian (Chernih et al. 2014). The idea that the level of differentiability is related to the conditioning of the gain matrix (Chernih et al. 2014) is supported by the observation that for the Wendland functions, the gain matrix for SGPR becomes numerically not-positive definite at a larger γ. The reverse happens for the Gaussian. For SGPR, the covariance function is differentiated twice (see Eqs. 8 and 10), reducing the level of differentiability of the resulting function for the Wendland functions, explaining the better conditioning of the matrix. Clearly, ϕ_G limits our choice of correlation length, whereas Wendland functions are always suitable.

4 Synthetic test cases

The performance of SGPR is assessed using two synthetic test cases of incompressible flow. It is compared with three other filters. The first does not enforce the divergence-free constraint, and works through a convolution of the data using a 3×3 smoothing kernel with equal weights ($\text{BOX}_{3\times3}$), a simple yet widely used approach to dealing with noisy data (Raffel et al. 1998). The two solenoidal filters we compare with are DCS (de Silva et al. 2013) and SWR (Schiavazzi et al. 2014). A detailed description of the specific implementations of these two filters in the present paper can be found in the Appendices 1 and 2.

The covariance function we use when applying SGPR is the Wendland function with smoothness C^4:

$$\phi_{d,2}(r) = (1 - r)_+^6 \left(\frac{35}{3} r^2 + 6r + 1 \right), \tag{13}$$

where $+$ indicates that $(1 - r)_+^6 = 0$ for $r \geq 1$. The choice for using this function is computational efficiency; it is the lowest-order Wendland function with continuous second-order velocity components, which are needed for calculating the pressure field, and as is concluded in Sect. 3.2, using lower-order covariance functions improves the conditioning of the gain matrix and as a result improves the convergence speed of the conjugate gradient method. Recall from Eq. 3 that the radial distance is scaled with a correlation length γ. The user can either specify its value a priori or its optimum value can be pursued by using an optimization algorithm. Choosing between the manual and

automatic approach is a trade-off between computational cost and accuracy. In this paper, we suggest to set the tuning parameter a priori. The reasoning is that in practice, an experimenter will already know what the flow looks like before applying SGPR. So the tuning parameter could be set equal to a representative flow feature, like the size of a vortex. In this section and the next, we will investigate how the reconstructed fields are influenced by the choice of the tuning parameter and then determine how the a priori choice compares with the actual optimum value.

We set $R = \lambda I$, where λ simply acts as a regularization parameter. Many different methods have been proposed in the literature to choose this parameter; however, there is no generally accepted approach. For an extensive review of available approaches, the reader is referred to Bauer and Lukas (2011). We chose $\lambda = 10^{-2}$ for reasons of computational efficiency; satisfactory convergence speeds were then obtained for all test cases when using the CG method.

To simulate PIV noise, we add *spatially correlated* Gaussian noise to the true velocity field. However, the error of one velocity component is assumed to be uncorrelated from the error of the other velocity components. This is typically the case in PIV. Therefore, noise will be added to each velocity component independently. For each velocity component, we express the covariance of the measurement uncertainty between point i and j as:

$$S^{ij} = \mathrm{Cov}\left[u_{\mathrm{PIV}}^i, u_{\mathrm{PIV}}^j\right] = \sigma^i \zeta^{ij} \sigma^j,$$

where σ^i and σ^j are the standard deviations of the measurement noise at points i and j, respectively. We take this standard deviation equal to a fraction of the local velocity magnitude. For point i:

$$\sigma^i = \alpha ||\mathbf{u}_{\mathrm{PIV}}^i||_2,$$

where we have taken $\alpha = 0.1$, i.e., the standard deviation is equal to 10 % of the local velocity magnitude. Therefore, the measurement uncertainty is spatially varying. This is roughly representative of the situation in a real PIV experiment, see e.g., Fig. 8. The correlation ζ^{ij} we use is based on results obtained by Wieneke and Sciacchitano (2015), who investigated the spatial correlation of measurement noise between PIV velocity vectors as a function of overlap and interrogation window size. We used the results for an overlap of 75 % and an interrogation window size of 32 voxels in each direction, which are the most frequently used settings for tomographic PIV (Scarano 2013). To eventually obtain noise with the desired spatial covariance structure, we compute the Cholesky decomposition of $S = LL'$, and multiply L with a vector containing *uncorrelated* noise from a standard Gaussian distribution. This noise is added to the analytical velocity fields.

The filters we will use work on the velocity field. To get the vorticity field, we take the curl. With SGPR we can calculate this field analytically (see Eq. 10). For the other filters we use finite differencing, as velocities are only available on grid nodes. To allow for a fair comparison, we also use finite differencing to compute curl when using SGPR.

For the first test case, we will also calculate the pressure field. The pressure field can be obtained from the velocity field by using the Navier–Stokes equations (van Oudheusden 2013):

$$\nabla p = -\rho \left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} \right) + \mu \nabla^2 \mathbf{u}. \tag{14}$$

By taking the divergence of Eq. 14, we obtain the pressure Poisson equation:

$$\nabla^2 p = \nabla \cdot \left\{ -\rho \left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} \right) + \mu \nabla^2 \mathbf{u} \right\}. \tag{15}$$

Neumann boundary conditions can be obtained from Eq. 14. All the derivatives are calculated using second-order accurate finite-difference schemes.

The metric we use to compare the filters is a noise reduction percentage (Q), defined as:

$$Q = \left(\frac{\varepsilon^{\mathrm{o}} - \varepsilon^{\mathrm{f}}}{\varepsilon^{\mathrm{o}}} \right) \times 100 \, \%, \tag{16}$$

where ε^{o} is the root-mean-square error of the measured (observed) field and ε^{f} is the root-mean-square error of the filtered field.

4.1 Taylor vortex

The Taylor vortex is defined in 2D (Panton 2013). The velocity field is solenoidal and in addition satisfies the Navier–Stokes equations. The tangential velocity is defined as:

$$u_\theta = \frac{H}{8\pi} \frac{r}{\nu t^2} \exp\left(-\frac{r^2}{4\nu t} \right), \tag{17}$$

where H is the amount of angular momentum in the vortex, r is the radial distance from the center of the vortex, ν is the kinematic viscosity and t is time. The vorticity is given by the following expression:

$$\omega = -\frac{H\left(r^2 - 4\nu t\right)}{16\pi \, \nu^2 t^3} \exp\left(-\frac{r^2}{4\nu t} \right). \tag{18}$$

The pressure is given by:

$$p = -\frac{\rho H^2}{64\pi^2 \nu t^3} \exp\left(-\frac{r^2}{2\nu t} \right) + p_\infty, \tag{19}$$

	BOX$_{3\times3}$ (%)	DCS (%)	SWR (%)	SGPR (%)
Velocity	14.9	29.6	35.8	59.2
Vorticity	26.3	0	15.8	70.4
Pressure	42.7	83.0	88.1	92.3

where ρ is the fluid density and p_∞ is the pressure at an infinite distance from the vortex center. For the present problem, we take $H = 1 \times 10^{-6}$ m^2, $\nu = 1 \times 10^{-6}$ m^2 s^{-1}, $\rho = 1000$ kg m^{-3}. The domain size is $-L_0 \leq x_i \leq L_0$, $i = 1, 2$, where $L_0 = 1 \times 10^{-3}$ m. We use a regular grid with 101 points in both directions, totaling 10,201 observation points. Using a regular grid allows us to solve the linear system efficiently by exploiting its Toeplitz structure. Samples are taken between $t = 0.05$ s and $t = 0.30$ s with a sampling frequency of 100 Hz. These same settings were also used by Charonko et al. (2010), who investigated the effect of a number of factors, among others measurement noise, on PIV-based pressure reconstruction.

For simplicity, we use the same correlation length in both directions. Since the whole domain contains one vortex, we set its value equal to the size of the domain: $\gamma = \gamma_1 = \gamma_2 = 2L_0$. Table 2 shows the noise reduction Q for the velocity magnitude, vorticity and pressure fields, averaged over the 26 time frames. Comparing SGPR with the other filters, we see that it provides the highest noise reduction for all fields. Comparing BOX$_{3\times3}$ with DCS and SWR, we see that BOX$_{3\times3}$ returns a more accurate vorticity field, but less accurate velocity and pressure fields. Notice that for DCS the noise in the vorticity field is not reduced, so the vorticity of the solenoidal field is equal to that of the observed field. This indicates that it is equivalent to the Helmholtz representation theorem. Comparing the three solenoidal filters, SGPR has the highest noise reduction, followed by SWR and then DCS.

To get a qualitative impression of the reconstructed fields, we refer to Fig. 2. First of all, notice that the noise we have added to represent the PIV noise is indeed spatially correlated. The fields returned by SGPR do indeed look better than the ones obtained with the other filters, especially when looking at the velocity and vorticity fields. The improvement in the pressure field is less noticeable, as verified from Table 2.

The correlation length we used was determined based on the fact that the field contains one vortex that encompasses the entire domain, so we set $\gamma / L_0 = 2$. Figure 3 shows how the noise reduction for the velocity, vorticity and pressure fields is influenced by the correlation length. The three graphs share the following characteristics: (1) There is an optimum noise reduction inside the interval; (2) small values of the correlation length lead to an abrupt decrease in noise reduction; (3) around the optimum, the graph is relatively flat. The optimum correlation length is around $\gamma / L_0 = 9$, resulting in noise reductions of 78, 93 and 93 % for velocity, vorticity and pressure, respectively. Though our choice of correlation length was therefore suboptimal, the a priori physical knowledge we used has still given us better noise reduction than the other filters. It is only when we would have chosen a length smaller than $\gamma / L_0 = 0.45$ that SWR would have been better. The fact that the graphs are relatively flat around the optimum is a desirable property as it implies robustness of SGPR. As long as the user makes an educated guess for the correlation length based on a priori physical knowledge, SGPR is relatively insensitive to the exact value chosen.

4.2 Vortex ring

The vortex ring has become a popular test case to numerically assess reconstruction methods in tomographic PIV (Elsinga et al. 2006; Novara and Scarano 2013). The displacement field used for it is a Gaussian, which does not result in a divergence-free velocity field. Therefore, we have used a different model proposed by Kaplanski et al. (2009). They derived analytical expressions for the velocity and vorticity fields by linearizing the vorticity equation for incompressible flows. Figure 4 shows an isosurface of the vorticity magnitude and the velocity field at the symmetry plane. The axial (u_a) and radial (u_r) velocities are defined as:

$$u_a = \pi \theta^2 \int_0^\infty \mu F(\mu, \eta) J_1(\theta \mu) J_0(\sigma \mu) \, d\mu, \tag{20}$$

$$u_r = -\pi \theta^2 \int_0^\infty \mu \tilde{F}(\mu, \eta) J_1(\theta \mu) J_1(\sigma \mu) \, d\mu, \tag{21}$$

where J_0 and J_1 are Bessel functions. The functions $F(\mu, \eta)$ and $\tilde{F}(\mu, \eta)$ are defined as follows:
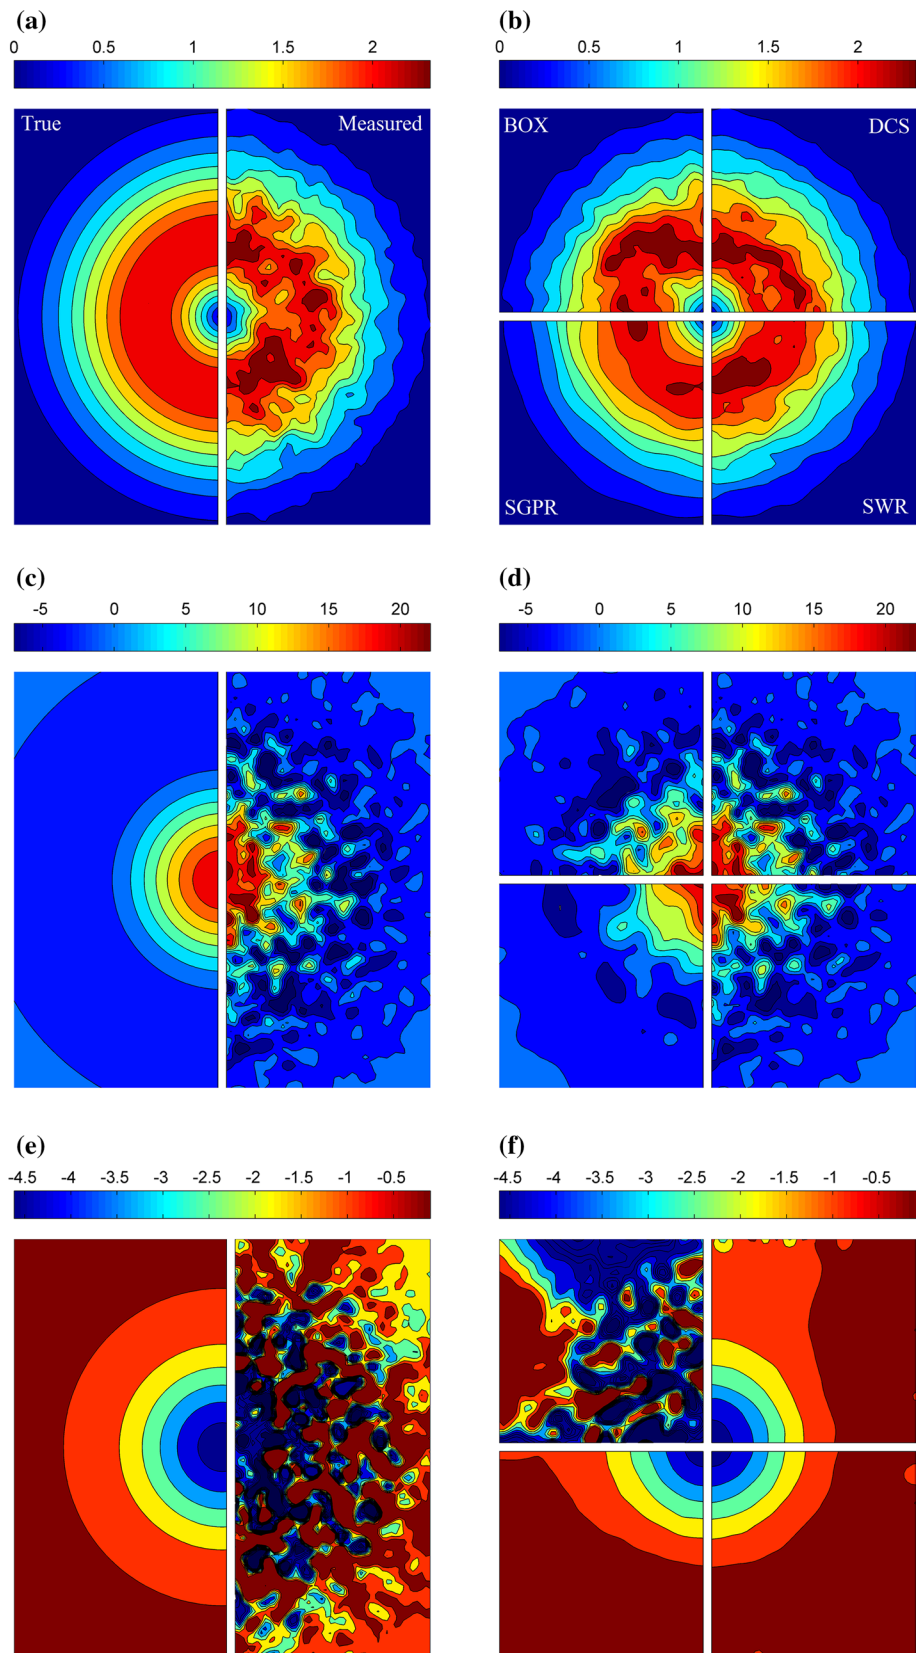
$$\begin{aligned} F(\mu, \eta) = \exp(\eta\mu)\mathrm{erfc}\left(\frac{\mu + \eta}{\sqrt{2}}\right) \\ + \exp(-\eta\mu)\mathrm{erfc}\left(\frac{\mu - \eta}{\sqrt{2}}\right), \end{aligned} \tag{22}$$

$$\begin{aligned} \tilde{F}(\mu, \eta) = \exp(\eta\mu)\mathrm{erfc}\left(\frac{\mu + \eta}{\sqrt{2}}\right) \\ - \exp(-\eta\mu)\mathrm{erfc}\left(\frac{\mu - \eta}{\sqrt{2}}\right). \end{aligned} \tag{23}$$

The parameters σ, η and θ are defined as:

$$\sigma = \frac{r}{l}, \quad \eta = \frac{x_1}{l}, \quad \theta = \frac{R_0}{l}.$$

Fig. 2 True (*left* side of **a**, **c**, **e**), measured (*right* side of **a**, **c**, **e**) and filtered (**b**, **d**, **f**) velocity magnitude (*top*), vorticity (*center*) and pressure (*bottom*) fields for the Taylor vortex. Filters (clockwise, starting at *top left*): BOX$_{3\times3}$, DCS, SWR, SGPR. The second frame is shown
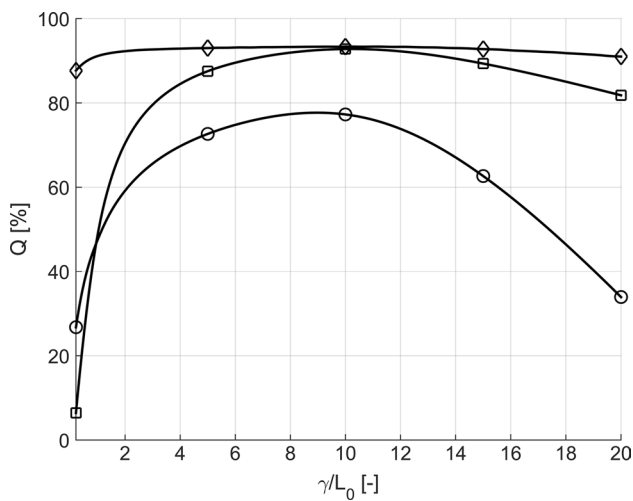
Fig. 3 Noise reduction in the velocity (*circle*), vorticity (*square*) and pressure (*diamond*) fields as a function of correlation length for the Taylor vortex using SGPR
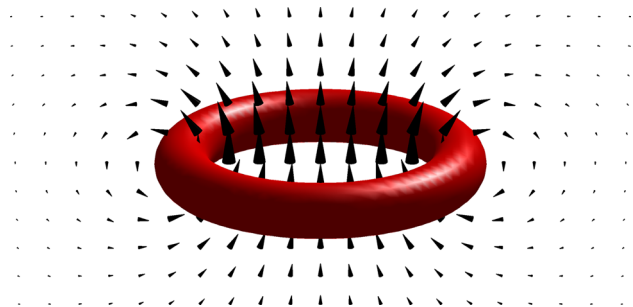
Fig. 4 Isosurface of vorticity magnitude at $2000\,(1/s)$ and velocity field at the symmetry plane for the vortex ring case

We take $R_0 = 0.0067$ and $l = 0.0031$. The vorticity is given by the following expression:

$$\omega = \frac{\Gamma_0 R_0}{\sqrt{2\pi}\, l^3} \exp\left[-\frac{1}{2}\left(\sigma^2 + \eta^2 + \theta^2\right)\right] I_1(\sigma\theta) \qquad (24)$$

where I_1 is the modified Bessel function. We take $\Gamma_0 = 4\pi R_0$. The domain size is $-2.5R_0 \leq x_1 \leq 2.5R_0$ and $-5R_0 \leq x_2, x_3 \leq 5R_0$. The coordinate direction 1 is taken as the axial direction. We take a regular grid with 21 data points in the 1-direction and 41 data points in the 2- and 3-directions. So there are 35,301 data points in total. We use a correlation length of $10R_0$ in each direction. So the support radius spans the complete vortex ring. Table 3 shows the noise reduction for the velocity and vorticity magnitude. Again, $\text{BOX}_{3\times3}$ results in the lowest noise reduction for the velocity field. However, as for the previous test case, it returns more accurate vorticity fields than DCS and SWR, where DCS again leaves the measured vorticity field unchanged. SGPR results in the highest noise

Table 3 Noise reduction Q of velocity magnitude and vorticity for the vortex ring

	$\text{BOX}_{3\times3}$	DCS (%)	SWR (%)	SGPR (%)
Velocity	16.6	17.6	27.5	27.4
Vorticity	18.5	0	15.1	20.7

reduction for the vorticity field, but for the velocity field SWR is marginally better. To get a qualitative impression of the methods, Fig. 5 compares the true vortex ring with the filtered ones. Figure 6 shows how the noise reduction is influenced by the choice of correlation length. The value we chose ($\gamma/R_0 = 10$) is close to the optimum at $\gamma/R_0 = 14$. At this value, the noise reduction in velocity and vorticity is 30.4 and 23.2 %, respectively. Similar to the Taylor vortex, the current graph has three characteristics, namely an optimum inside the interval, a large decrease in noise reduction for small correlation lengths and a relatively flat curve around the optimum.

5 Application to experimental data

Two experimental data sets are considered: The first is of a circular jet in water (Violato and Scarano 2011); the second is of a fully developed turbulent boundary layer over a flat plate in air (Pröbsting et al. 2013). Both data sets were obtained using tomographic PIV. We will not consider the non-solenoidal $\text{BOX}_{3\times3}$ filter, since such filtering steps have already been applied to obtain the velocity fields from the particle images, especially due to the effect of the interrogation windows.

To assess the reconstruction accuracy of the solenoidal filters, we use different approaches for the two test cases: For the circular jet in water, we use a posteriori estimates of random uncertainty for the measurements in terms of the velocity field and investigate to what extent the filtered velocities are inside the uncertainty bounds; for the turbulent boundary layer in air, we reconstruct the pressure field from the velocity data and compare with microphone measurements.

5.1 Circular jet in water

Figure 7 shows the circular jet in water. The jet velocity at the nozzle exit is 0.5 m/s, so assuming incompressible flow is an excellent approximation. The jet periodically creates vortex rings, as can clearly be seen from Fig. 7. The data set consists of 289,328 vectors and is defined on a regular grid, with 107 vectors in the jet direction and 52 vectors in the cross directions. Since we are considering experimental

Fig. 5 *Top view* of the vortex ring and contours of the out-of-plane velocity. True (*left* side of **a**), observed (*right* side of **a**) and filtered **b** fields
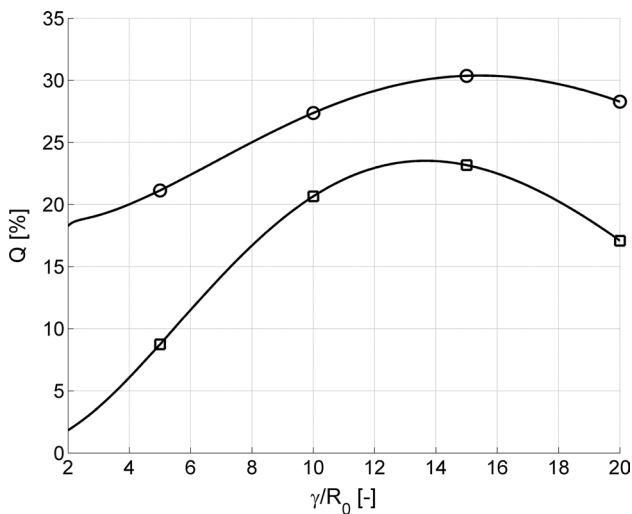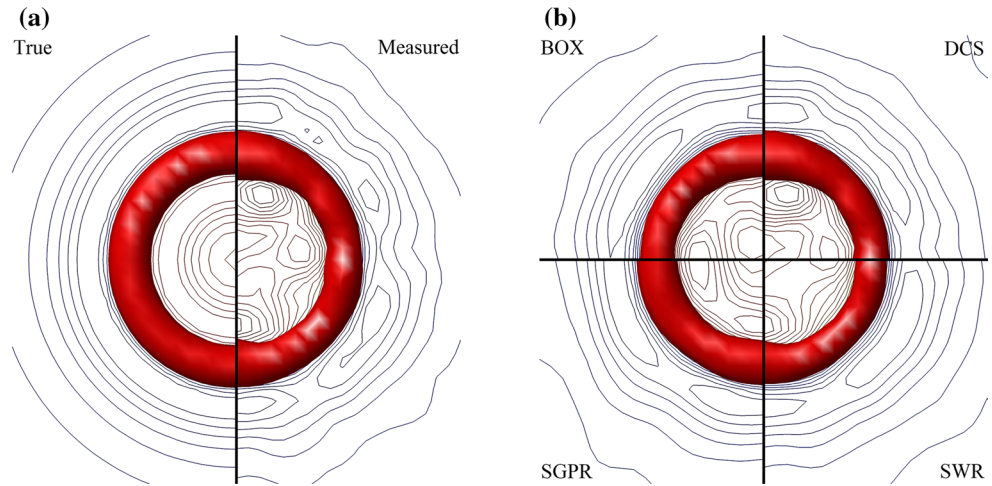

Fig. 6 Noise reduction in the velocity (*circle*) and vorticity (*square*) fields as a function of correlation length for the vortex ring using SGPR

data, there is no perfect reference solution. To make a quantitative performance assessment of the solenoidal filters, we use a recent method that estimates a posteriori random uncertainty for measurements locally (Sciacchitano 2015). It is also applicable for tomographic PIV. It expresses the distribution of error for each vector component in terms of standard deviation (1 sigma). Figure 8 shows isosurfaces of the standard deviation at a velocity magnitude of 0.035 (m/s). Clearly, the measurement uncertainty varies in space. Also notice how the uncertainty is localized inside the jet, as opposed to outside where the flow is practically at rest. In particular, we can identify ring shapes (compare with the vortex rings in Fig. 7). These are regions with strong velocity gradients, so it is reasonable to expect larger uncertainties there (Sciacchitano et al. 2013). When using SGPR, we consider two scenarios. In the first, we do not include the

a posteriori measurement uncertainty in the reconstruction. In the second scenario, we take it into account in R.

The covariance function we use is given by Eq. 13. The correlation length we use spans 15 vectors in each direction, which is approximately the diameter of a vortex ring, i.e., $d_{\mathrm{vortex}} = 15h$. To investigate the influence of the correlation length on the reconstructed velocity field, we under sample the data set by halving the measurement resolution in each direction. The data points in between are taken as validation, which we then use to calculate the root-mean-square error. Figure 9 shows the result. The behavior is very similar to what we found for the synthetic data sets in the previous section: (1) There is an optimum correlation length inside the interval; (2) small values of the correlation length lead to an abrupt increase in root-mean-square error; (3) around the optimum, the graph is relatively flat . Again, we emphasize that the latter characteristic is a desirable one from a robustness point of view. It turns out that our choice of correlation length $\gamma/d_{\mathrm{vortex}} = 1$ was a very good one. Not only are we in the flat region, we are at the beginning of it. This is advantageous since a smaller correlation length leads to a better conditioned gain matrix, resulting in faster convergence. Note that we have only under sampled to understand the relation between reconstruction accuracy and correlation length. In the following, we use the full data set.

Figure 10a shows isosurfaces of the velocity divergence at 15 (1/s). Before removing the spurious divergence, we first quantify how far from divergence-free the data set is. The following metric is due to Zhang et al. (1997):

$$\delta = \frac{(\partial u_1/\partial x_1 + \partial u_2/\partial x_2 + \partial u_3/\partial x_3)^2}{(\partial u_1/\partial x_1)^2 + (\partial u_2/\partial x_2)^2 + (\partial u_3/\partial x_3)^2}. \tag{25}$$

The mean of δ will lie between 0 and 1. When $\bar{\delta} = 0$, the velocity field is exactly divergence-free. When $\bar{\delta} = 1$, the velocity components are independent, random variables.

Fig. 7 Circular jet in water, obtained with tomographic PIV. Velocity vector slice in the axial plane. Isosurfaces of vorticity magnitude at $\omega = 220$ (1/s)
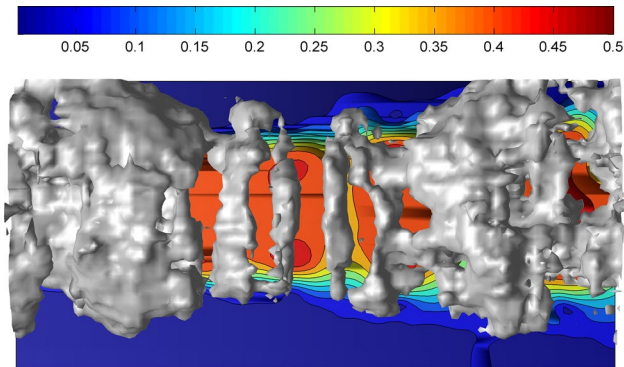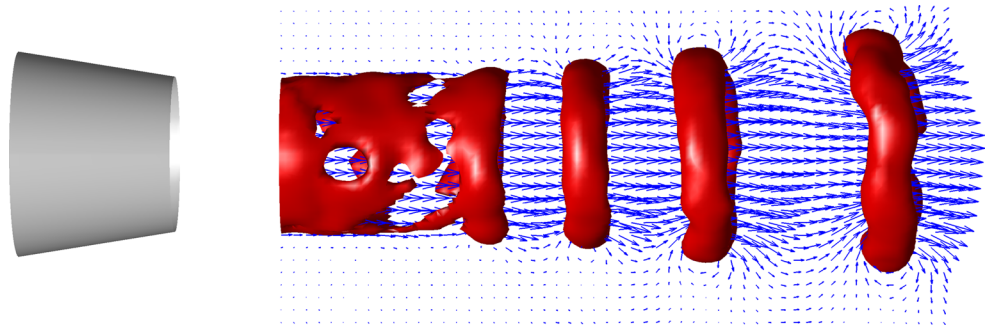

Fig. 8 Velocity magnitude slice in the axial plane. Isosurfaces of measurement uncertainty in terms of velocity magnitude standard deviation at 0.035 (m/s)
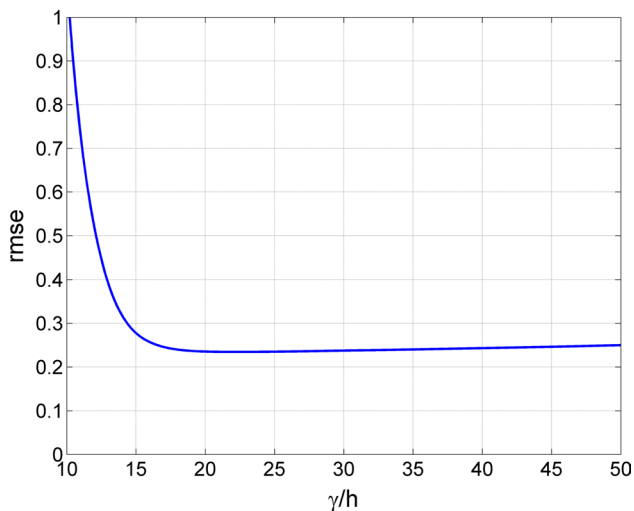

Fig. 9 Root-mean-square error as a function of correlation length for the circular jet in water

Fig. 10 Isosurfaces of velocity divergence at 15 (1/s) of the original data set (**a**) and after applying SGPR (**b**)

Zhang et al. (1997) showed that smoothing of the data can significantly reduce this value. For a turbulent flow measurement in a square duct using hybrid holographic PIV, they showed how using a Gaussian filter decreased $\bar{\delta}$ from 0.74 to 0.50. For the data set used in the present paper, we

found that $\bar{\delta} = 0.66$. Figure 10b shows the isosurfaces of velocity divergence after applying SGPR. The mean of the normalized divergence is $\bar{\delta} = 0.09$. Note that even though SGPR defines an analytically divergence-free field, we see some spurious divergence. This is because it was calculated using finite differencing to compare with the original data set. Therefore, the spurious divergence we observe for SGPR is completely due to truncation error. So the spurious divergence observed in the data set is for a large part not due to truncation error, but other sources that arise from measurements, like random measurement errors. The reason that we do not show the velocity divergence after applying DCS or SWR is because they are defined to return divergence-free fields on a finite-difference and

finite-volume level, respectively. So, the divergence calculated with these methods at the data points is zero.

We quantify the performance of the filters by using the following metric:

$$\xi_i(\mathbf{x}) = \frac{\left| u_{\mathrm{sol},i}(\mathbf{x}) - u_{\mathrm{exp},i}(\mathbf{x}) \right|}{\sigma_i(\mathbf{x})}, \quad i = 1, 2, 3 \qquad (26)$$

where $u_{\mathrm{sol},i}(\mathbf{x})$ is the solenoidal velocity field in the i-direction, obtained with DCS, SWR or SGPR and $u_{\mathrm{exp},i}(\mathbf{x})$ is the experimental velocity in the i-direction. The a posteriori standard deviation of the random uncertainty in the i-direction is given by $\sigma_i(\mathbf{x})$. Ideally, all reconstructed data points should have $\xi_i \le 3$, i.e., the error should be within 3 standard deviations of the measurement uncertainty (99.7 % certainty). Figure 11 shows the cumulative distribution of ξ_1. The plots for the other two directions look very similar. The results for DCS and SWR are practically identical. When the a posteriori measurement uncertainty information is not included (1), SGPR is only slightly better than the other filters. However, when this information is included in the reconstruction (2), the method is clearly better in reconstructing a solenoidal velocity field within the measurement uncertainty. At $\xi_i = 3$, $F(\xi_i)$ is 1, whereas for DCS and SWR the value is around 0.9, meaning that 10 % of the data points have a deviation from the measurements greater than 3 standard deviations. With regard to the average ξ_1, DCS and SWR are equal to 2.0. For SGPR, it is 1.8 when no measurement uncertainty is included, but it reduces to 0.3 when this information is incorporated. In conclusion, when including local measurement uncertainty, SGPR is able to return a divergence-free velocity field that follows the measured velocity field more faithfully.
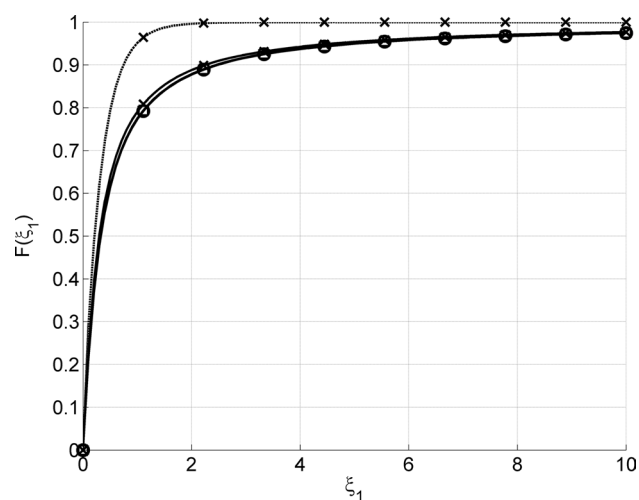
5.2 Turbulent flat plate boundary layer in air

For this data set, the free stream velocity is 10 m/s, so incompressible flow is again an excellent approximation. The data set consists of 185,484 vectors and is defined on a regular grid, with 58 vectors in streamwise (x)-direction, 123 vectors in the spanwise (z)-direction and 26 vectors in the wall-normal (y) direction. Since the particle images were not available to us, we were unable to use the approach by Sciacchitano (2015) to estimate the measurement uncertainty, as was done in the first experimental test case. Therefore, for SGPR the observation error covariance matrix R is set to the identity matrix multiplied with $\lambda = 10^{-2}$. We again use Eq. 13 as the covariance function. Recall that for the circular jet in water, we chose the correlation length equal to the diameter of the vortex ring. The present turbulent boundary layer does not have such a distinguishable feature. One can suggest to use the boundary layer thickness for this and set the correlation length equal to it. However, the boundary layer thickness is approximately 9.4 mm and the height of the measurement volume is 4.2 mm, so the measurement volume is completely immersed in the boundary layer. For this reason, we simply chose the correlation length equal to the height of the measurement volume H. Indeed, as was explained before, larger correlation length results in a more ill-conditioned system matrix, which increases the computational time. To investigate the influence of the correlation length on the reconstructed velocity field, we again under sample the data set by halving the measurement resolution in each direction. The data points in between are taken as validation, which we then use to calculate the root-mean-square error. Figure 12 shows the result. We observe that the
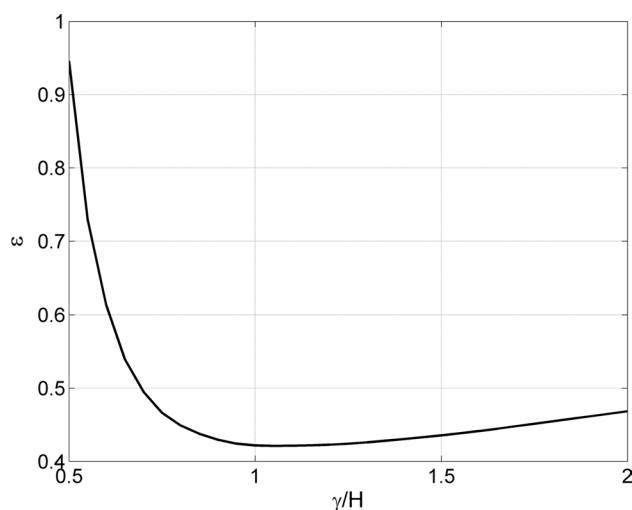

Fig. 11 Cumulative distribution of ξ_1 for the circular jet in water using DCS (*circle*), SWR (*square*) and SGPR (1) (*times, solid*) and SGPR (2) (*times, dashed*)

Fig. 12 Root-mean-square error as a function of correlation length for the turbulent boundary layer in air

chosen correlation length of $\gamma/H = 1$ is close to the optimum. Before applying SGPR, $\bar{\delta} = 0.74$. After application of this filter, it reduces to $\bar{\delta} = 0.03$.

Contrary to the previous experimental test case, no a posteriori estimates of random uncertainty for the velocity measurements are available. Instead, to quantify the reconstruction accuracy of the solenoidal filters, we derive pressure fields from the time-resolved velocity field measurements and compare the results with a pinhole microphone located on the wall. The boundary conditions imposed are the same as discussed in Pröbsting et al. (2013), so Neumann boundary conditions at all sides except the top side closest to the free stream, where the pressure is prescribed (Dirichlet boundary condition). Three approaches are followed to evaluate the acceleration, required for the pressure calculation: (1) the Eulerian scheme (eul), where the material acceleration is evaluated with respect to a stationary reference frame; (2) the "standard" Lagrangian scheme (lag), where the material acceleration is calculated by following a fluid particle. The specific method used is that mentioned by van Oudheusden (2013), i.e., a first-order reconstruction of the particle track with a central difference scheme for the material acceleration at the next and previous time instances; (3) a least-squares regression of the fluid parcel's velocity using a first-order polynomial basis (lsr_N), as proposed by Pröbsting et al. (2013). Here, N is the number of velocity fields over which the particle trajectory is followed. The two latter approaches are both Lagrangian methods. According to Violato et al. (2011), Lagrangian methods are more accurate for convection dominated flows, like boundary layer flows. For the synthetic test case of the (non-convecting) Taylor vortex (Sect. 4.1), we only used the Eulerian scheme. To determine the microphone location on the wall (the particle images were not available), we cross-correlated the time series of the microphone measurement with the PIV-derived pressure fields. Figure 13 shows the cross-correlation coefficient taken at a region where we believed the microphone was located. We clearly see a peak and this is where the microphone is most likely located.

Following Pröbsting et al. (2013), we show the pressure time series for a subset of the data in Fig. 14. Both the microphone and PIV-derived pressure signals were band pass-filtered over the range 300 Hz to 3 kHz. We have not shown the signals obtained with the other solenoidal filters to avoid obscuring the plots; however, they are very similar to SGPR. As expected, the worst reconstruction is obtained with the Eulerian scheme. Improvements are obtained by switching to the Lagrangian scheme. The best results are obtained when using the least-squares regressor, where 9 velocity fields were used for the particle trajectory. Comparing the results before and after applying SGPR, we clearly see that SGPR improves the reconstruction. The improvement is most noticeable for the Eulerian
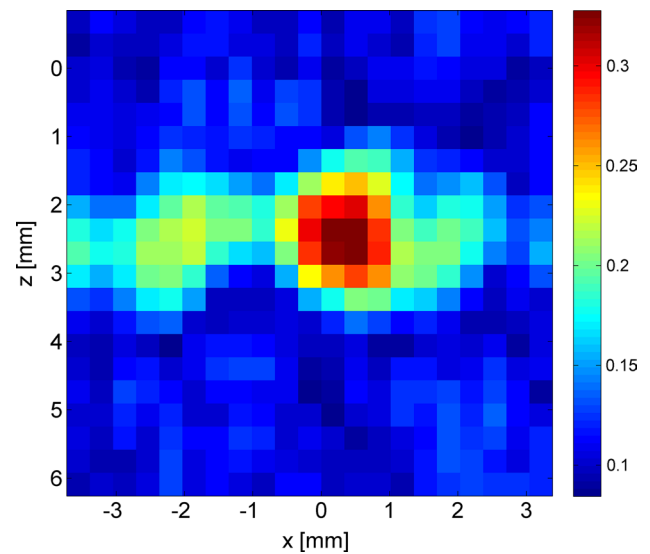

Fig. 13 Cross-correlation between microphone pressure time series and PIV-reconstructed pressure time series, as a function of location on the wall
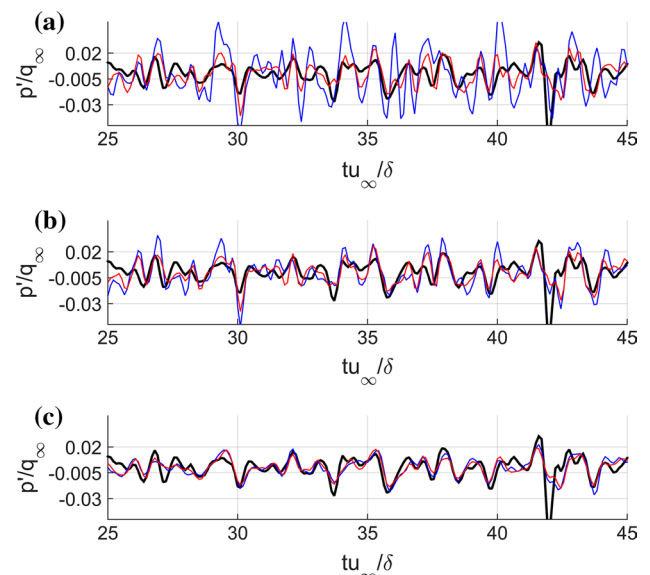

Fig. 14 Pressure time series obtained with the microphone (*black*), tomo PIV with the original data (*blue*) and SGPR (*red*) using the Eulerian (**a**), standard Lagrangian (**b**) and least-squares regression (lsr$_9$) (**c**) approaches. δ is the boundary layer thickness, u_∞ is the free stream velocity, q_∞ is the free stream dynamic pressure and p' the pressure fluctuation

scheme and the least for the least-squares regressor with 9 velocity fields, in which case the random errors seem to be sufficiently suppressed, leaving out mainly bias errors. However, we emphasize that using many velocity fields to reconstruct a particle trajectory is in principle undesirable since this results in a larger number of particles to fall

Table 4 Cross-correlation coefficient achieved with the original data (org) and the solenoidal filters

	eul	lag	lsr$_3$	lsr$_5$	lsr$_7$	lsr$_9$
org	0.38	0.46	0.46	0.55	0.60	0.62
DCS	0.51	0.55	0.56	0.62	0.63	0.65
SWR	0.51	0.54	0.55	0.61	0.62	0.63
SGPR	0.51	0.55	0.56	0.62	0.63	0.65

outside the measurement domain, potentially cropping a large part of the reconstructed pressure field. Therefore, it is indeed an important advantage that the solenoidal filters result in a significant improvement when using the Eulerian and standard Lagrangian schemes. For completeness, Table 4 shows the cross-correlation coefficients achieved using the various solenoidal filters and the various acceleration calculation schemes. As stated previously, the solenoidal filters return similar results and they always result in an improved pressure signal. This is independent of which scheme is used to calculate the acceleration, though the gain becomes less noticeable as more velocity fields are used to reconstruct a particle trajectory using lsr.

6 Conclusions

We investigated the use of an analytically solenoidal filter based on Gaussian process regression (SGPR) to remove spurious divergence from volumetric velocity measurements. We formulated the filter from a Bayesian perspective. Measurement uncertainty can therefore be included naturally. To allow filtering large data sets, we have exploited the Toeplitz structure of the resulting gain matrix to speed up matrix-vector multiplications in the CG method. This can be done when the measurements are on a (near) regular grid. We used two synthetic test cases to assess the accuracy of the filtered velocity and derived vorticity and pressure fields and compared with conventional (non-physics-based) filtering and two other recently proposed solenoidal filters that have been applied to volumetric velocity measurements as well. We have used a realistic model of PIV noise with *correlated* random errors. We found that SGPR has better noise reduction properties. The advantage of the other solenoidal filters is that they do not require a tuning parameter. SGPR on the other hand does, but we found that as long as the user makes use of his a priori physical knowledge of the field of interest, the results are quite insensitive to the exact value chosen. This demonstrates that the method is robust to its tuning parameters. The first experimental data set showed that by including the measurement uncertainty, SGPR is able to reconstruct a solenoidal velocity field that more faithfully follows the measurements. The second experimental data set showed that all three solenoidal filters improve the reconstructed pressure signal, independent

of which scheme is used to evaluate flow acceleration. Feel free to contact the second author to obtain the Matlab codes of the three solenoidal filters.

Appendix 1: DCS implementation

The divergence correction scheme (DCS), as proposed by de Silva et al. (2013), works as follows. It defines the following optimization problem:

$$\min_{\mathbf{u}_{\text{sol}}} \frac{1}{m} \sum_{i=1}^{m} \left[\sum_{j=1}^{3} \left(u_{\text{sol},j}^{i} - u_{\text{PIV},j}^{i} \right)^2 \right] \quad \text{subject to} \tag{27}$$

$$\nabla \cdot \mathbf{u}_{\text{sol}} = 0 \tag{28}$$

The divergence-free condition in Eq. 28 is approximated using a second-order central difference scheme on the Cartesian measurement grid. To solve the optimization problem, de Silva et al. (2013) use MATLAB's constrained nonlinear multivariable solver *fmincon*. However, a closer look at the defined optimization problem reveals that it can be restated as a quadratic programming problem with linear equality constraints (Wright and Nocedal 1999), therefore avoiding the need for Matlab's Optimization Toolbox:

$$\min_{\mathbf{u}_{\text{sol}}} \frac{1}{2}\mathbf{u}'_{\text{sol}}Q\mathbf{u}_{\text{sol}} + \mathbf{d}'\mathbf{u}_{\text{sol}} \quad \text{subject to} \tag{29}$$

$$D\mathbf{u}_{\text{sol}} = 0, \tag{30}$$

where $\mathbf{u}_{\text{sol}} = \left(\mathbf{u}_{\text{sol},1} \ \mathbf{u}_{\text{sol},2} \ \mathbf{u}_{\text{sol},3} \right)'$, $Q = (2/m)\, I$ (I is the identity matrix), $\mathbf{d} = -(2/m)\left(\mathbf{u}_{\text{PIV},1} \ \mathbf{u}_{\text{PIV},2} \ \mathbf{u}_{\text{PIV},3} \right)'$, and D is the discrete divergence operator. The solution to the

optimization problem can be obtained by solving the following linear system of equations:

$$\begin{bmatrix} Q & D' \\ D & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u}_{\text{sol}} \\ \boldsymbol{\lambda} \end{bmatrix} = \begin{bmatrix} -\mathbf{d} \\ 0 \end{bmatrix}, \tag{31}$$

where $\boldsymbol{\lambda}$ are Lagrange multipliers. The system matrix is sparse indefinite and can be solved using SYMMLQ or MINRES (Paige and Saunders 1975).

Our reformulation of DCS as a quadratic programming problem with linear equality constraints reveals that it is equivalent to the Helmholtz representation theorem. To this end, we extract the first row from Eq. 31:

$$Q\mathbf{u}_{\text{sol}} + D'\boldsymbol{\lambda} = -\mathbf{d}, \tag{32}$$

which can be rewritten to

$$\mathbf{u}_{\text{sol}} = \mathbf{u}_{\text{PIV}} - D'\boldsymbol{\lambda}. \tag{33}$$

The transpose of the divergence operator (D') is the nabla operator or gradient G (Ma et al. 2013):

$$\mathbf{u}_{\text{sol}} = \mathbf{u}_{\text{PIV}} - G\boldsymbol{\lambda}. \tag{34}$$

Taking the curl on both sides of Eq. 34 and recalling the vector identity that the curl of the gradient operator is zero, we find that the vorticity of the solenoidal velocity field is equal to the vorticity of the measured field. Indeed, this is equivalent to the Helmholtz representation theorem. This observation was also backed up by the synthetic test cases, where we found that DCS did not result in a noise reduction in the vorticity field.

Appendix 2: SWR implementation

Solenoidal waveform reconstruction (SWR), as proposed by Schiavazzi et al. (2014), works by first defining a Voronoi tessellation of the measurement grid. Then, the observed velocities $\mathbf{u}_{\text{PIV}}$ are converted to face fluxes $\mathbf{g}$. Vortices with unknown strengths $\boldsymbol{\alpha}$ are defined around all edges of the grid (in 2D, vortices are defined around nodes). This ensures that the reconstructed velocity field, which is a sum of these vortices, will be solenoidal. The vortex strengths are found by solving the system $W\boldsymbol{\alpha} = \mathbf{g}$ in the least-squares sense, where the sparse matrix W contains the influence of the solenoidal waveforms in terms of the flux they generate. Schiavazzi et al. (2014) propose a sequential matching pursuit algorithm to solve the underdetermined system. We use LSQR (Paige and Saunders 1982), a conjugate-gradient-type method for sparse least-squares systems. Once the vortex strengths have been found, the fluxes $\mathbf{g}^*$ are used to reconstruct the divergence-free velocities $\mathbf{u}_{\text{sol}}$.

References

Ammar GS, Gragg WB (1988) Superfast solution of real positive definite Toeplitz systems. SIAM J Matrix Anal Appl 9(1):61–76

Azijli I, Sciacchitano A, Ragni D, Palha A, Dwight RP (2015) *A posteriori* uncertainty quantification of PIV-derived pressure fields. In: 11th international symposium on PIV-PIV15

Battle G, Federbush P (1993) Divergence-free vector wavelets. Mich Math J 40(1):81–195

Bauer F, Lukas MA (2011) Comparing parameter choice methods for regularization of ill-posed problems. Math Comput Simul 81(9):1795–1841

Billings SD, Beatson RK, Newsam GN (2002) Interpolation of geophysical data using continuous global surfaces. Geophysics 67(6):1810–1822

Brücker C (1995) Digital-particle-image-velocimetry (DPIV) in a scanning light-sheet: 3D starting flow around a short cylinder. Exp Fluids 19(4):255–263

Busch J, Giese D, Wissmann L, Kozerke S (2013) Reconstruction of divergence-free velocity fields from cine 3D phase-contrast flow measurements. Magn Reson Med 69(1):200–210

Capizzano SS (2002) Matrix algebra preconditioners for multilevel Toeplitz matrices are not superlinear. Linear Algebra Appl 343:303–319

Chan RH, Jin XQ (2007) An introduction to iterative Toeplitz solvers, vol 5. SIAM, Philadelphia

Charonko JJ, Vlachos PP (2013) Estimation of uncertainty bounds for individual particle image velocimetry measurements from cross-correlation peak ratio. Meas Sci Technol 24(6):065,301

Charonko JJ, King CV, Smith BL, Vlachos PP (2010) Assessment of pressure field calculations from particle image velocimetry measurements. Meas Sci Technol 21(10):105401. doi:10.1088/0957-0233/21/10/105401

Chernih A, Sloan I, Womersley R (2014) Wendland functions with increasing smoothness converge to a Gaussian. Adv Comput Math 40(1):185–200

Darve E (2000) The fast multipole method: numerical implementation. J Comput Phys 160(1):195–240

Davis GJ, Morris MD (1997) Six factors which affect the condition number of matrices associated with kriging. Math Geol 29(5):669–683

Deriaz E, Perrier V (2006) Divergence-free and curl-free wavelets in two dimensions and three dimensions: application to turbulent flows. J Turbul 7(3):1468–5248. doi:10.1080/14685240500260547

de Silva CM, Philip J, Marusic I (2013) Minimization of divergence error in volumetric velocity measurements and implications for turbulence statistics. Exp Fluids 54(7):1–17

de Baar JH, Percin M, Dwight RP, van Oudheusden BW, Bijl H (2014) Kriging regression of PIV data using a local error estimate. Exp Fluids 55(1):1–13

Elkins CJ, Markl M, Pelc N, Eaton JK (2003) 4D magnetic resonance velocimetry for mean velocity measurements in complex turbulent flows. Exp Fluids 34(4):494–503

Elsinga GE, Scarano F, van Oudheusden BWBW (2006) Tomographic particle image velocimetry. Exp Fluids 41(6):933–947

Greengard L, Rokhlin V (1987) A fast algorithm for particle simulations. J Comput Phys 73(2):325–348

Gunes H, Rist U (2008) On the use of Kriging for enhanced data reconstruction in a separated transitional flat-plate boundary layer. Phys Fluids 20:104109. doi:10.1063/1.3003069

Hinsch KD (2002) Holographic particle image velocimetry. Meas Sci Technol 13(7):R61–R72

Inggs MR, Lord RT (1996) Interpolating satellite derived wind field data using Ordinary Kriging, with application to the nadir gap. IEEE Trans Geosci Remote 34(1):250–256

Jones DR, Schonlau M, Welch WJ (1998) Efficient global optimization of expensive black-box functions. J Glob Optim 13(4):455–492

Kailath T, Kung SY, Morf M (1979) Displacement ranks of a matrix. Am Math Soc B 1(5):769–773

Kaplanski F, Sazhin SS, Fukumoto Y, Heikal SBM (2009) A generalized vortex ring model. J Fluid Mech 622:233–258

Ko J, Kurdila AJ, Rediniotis OK (2000) Divergence-free bases and multiresolution methods for reduced-order flow modeling. AIAA J 38(12):2219–2232

Lee SL, Huntbatch A, Yang GZ (2008) Contractile analysis with kriging based on MR myocardial velocity imaging. In: Metaxas D, Axel L, Fichtinger G, Székely G (eds) Medical image computing and computer-assisted intervention-MICCAI 2008. Lecture notes in computer science, vol 5241. Springer, Berlin, pp 892–899

Liburdy JA, Young EF (1992) Processing of three-dimensional particle tracking velocimetry data. Opt Laser Eng 17(3):209–227

Lowitzsch S (2005) A density theorem for matrix-valued radial basis functions. Numer Algorithms 39(1–3):253–256

Ma Z, Chew W, Jiang L (2013) A novel fast solver for poisson's equation with neumann boundary condition. Prog Electromagn Res 136:195–209

Maas HG, Gruen A, Papantoniou D (1993) Particle tracking velocimetry in three-dimensional flows. Exp Fluids 15(2):133–146

Memarsadeghi N, Raykar VC, Duraiswami R, Mount DM (2008) Efficient Kriging via fast matrix-vector products. In: Aerospace conference, 2008 IEEE. IEEE, pp 1–7

Narcowich FJ, Ward JD (1994) Generalized Hermite interpolation via matrix-valued conditionally positive definite functions. Math Comput 63(208):661–687

Ng MK, Pan J (2010) Approximate inverse circulant-plus-diagonal preconditioners for Toeplitz-plus-diagonal matrices. SIAM J Sci Comput 32(3):1442–1464

Novara M, Scarano F (2013) A particle-tracking approach for accurate material derivative measurements with tomographic piv. Exp Fluids 54(8):1–12

Paige CC, Saunders MA (1975) Solution of sparse indefinite systems of linear equations. SIAM J Numer Anal 12(4):617–629

Paige CC, Saunders MA (1982) LSQR: an algorithm for sparse linear equations and sparse least squares. ACM Trans Math Softw 8(1):43–71

Panton R (2013) Incompressible flow. Wiley, New York

Pröbsting S, Scarano F, Bernardini M, Pirozzoli S (2013) On the estimation of wall pressure coherence using time-resolved tomographic PIV. Exp Fluids 54(7):1–15

Raffel M, Willert CE, Kompenhans J (1998) Particle image velocimetry: a practical guide; with 24 tables. Springer, Berlin

Rasmussen C, Williams C (2006) Gaussian processes for machine learning, vol 1. MIT, Cambridge

Sadati M, Luap C, Kröger M, Öttinger HC (2011) Hard vs soft constraints in the full field reconstruction of incompressible flow kinematics from noisy scattered velocimetry data. J Rheol 55(6):1187–1203

Scarano F (2013) Tomographic PIV: principles and practice. Meas Sci Technol 24(1):1–28

Schaback R (1995) Error estimates and condition numbers for radial basis function interpolation. Adv Comput Math 3(3):251–264

Scheuerer M, Schlather M (2012) Covariance models for divergence-free and curl-free random vector fields. Stoch Models 28(3):433–451

Schiavazzi D, Coletti F, Iaccarino G, Eaton JK (2014) A matching pursuit approach to solenoidal filtering of three-dimensional velocity measurements. J Comput Phys 263:206–221

Schneiders JF, Dwight RP, Scarano F (2014) Vortex-in-cell method for time-supersampling of PIV data. Exp Fluids 55(3)

Schräder D, Wendland H (2011) A high-order, analytically divergence-free discretization method for Darcys problem. Math Comput 80(273):263–277

Sciacchitano A (2015) *A posteriori* uncertainty quantification for tomographic PIV data. In: 11th international symposium on PIV-PIV15

Sciacchitano A, Dwight RP, Scarano F (2012) Navier–Stokes simulations in gappy PIV data. Exp Fluids 53(5):1421–1435

Sciacchitano A, Wieneke B, Scarano F (2013) PIV uncertainty quantification by image matching. Meas Sci Technol 24(4):045,302

Segel L (2007) Mathematics applied to continuum mechanics. SIAM, Philadelphia

Song SM, Napel S, Glover GH, Pelc NJ (1993) Noise reduction in three-dimensional phase-contrast MR velocity measurementsl. J Magn Reson Imaging 3(4):587–596

Suzuki T, Ji H, Yamamoto F (2009) Unsteady PTV velocity field past an airfoil solved with DNS: part 1. Algorithm of hybrid simulation and hybrid velocity field at Re $\approx 10^3$. Exp Fluids 47(6):957–976

Timmins BH, Wilson BW, Smith BL, Vlachos PP (2012) A method for automatic estimation of instantaneous local uncertainty in particle image velocimetry measurements. Exp Fluids 53(4):1133–1147

Urban K (1996) Using divergence free wavelets for the numerical solution of the Stokes problem. In: Axelsson O, Polman B (eds) Proceedings of the conference on algebraic multilevel iteration methods, Nijmegen, pp 259–278

van Oudheusden BW (2013) PIV-based pressure measurement. Meas Sci Technol 24(3):1–32

Vennell R, Beatson R (2009) A divergence-free spatial interpolator for large sparse velocity data sets. J Geophys Res-Oceans 114:C10024. doi:10.1029/2008JC004973

Violato D, Scarano F (2011) Three-dimensional evolution of flow structures in transitional circular and chevron jets. Phys Fluids 23(12):124,104

Violato D, Moore P, Scarano F (2011) Lagrangian and Eulerian pressure field evaluation of rod-airfoil flow from time-resolved tomographic PIV. Exp Fluids 50(4):1057–1070

Wendland H (2005) Scattered data approximation, vol 17. Cambridge University Press, Cambridge

Wendland H (2009) Divergence-free kernel methods for approximating the Stokes problem. SIAM J Numer Anal 47(4):3158–3179

Wieneke B, Prevost R (2014) DIC uncertainty estimation from statistical analysis of correlation values. In: Advancement of optical methods in experimental mechanics, vol 3. Springer, Berlin, pp 125–136

Wieneke B, Sciacchitano A (2015) PIV uncertainty propagation. In: 11th international symposium on PIV-PIV15

Wikle CK, Berliner M (2007) A Bayesian tutorial for data assimilation. Phys D 230(1):1–16

Worth N (2012) Measurement of three-dimensional coherent fluid structure in high Reynolds number turbulent boundary layers. PhD thesis, University of Cambridge

Wright SJ, Nocedal J (1999) Numerical optimization, vol 2. Springer, New York

Yang GZ, Kilner PJ, Firmin DN, Underwood SR, Burger P, Longmore DB (1993) 3D cine velocity reconstruction using the method of convex projections. In: Computers in cardiology 1993, Proceedings, IEEE, pp 361–364

Zhang J, Tao B, Katz J (1997) Turbulent flow measurement in a square duct with hybrid holographic PIV. Exp Fluids 23(5):373–381