

Sparsifying the resolvent forcing mode via gradient-based optimisation

Calum S. Skene^{1,†,‡}, Chi-An Yeh^{1,§}, Peter J. Schmid^{2,¶} and Kunihiko Taira¹

¹Department of Mechanical and Aerospace Engineering, University of California, Los Angeles, CA 90095, USA

²Department of Mathematics, Imperial College London, London SW7 2AZ, UK

(Received 28 September 2021; revised 6 March 2022; accepted 31 May 2022)

We consider the use of sparsity-promoting norms in obtaining localised forcing structures from resolvent analysis. By formulating the optimal forcing problem as a Riemannian optimisation, we are able to maximise cost functionals whilst maintaining a unit-energy forcing. Taking the cost functional to be the energy norm of the driven response results in a traditional resolvent analysis and is solvable by a singular value decomposition (SVD). By modifying this cost functional with the L_1 -norm, we target spatially localised structures that provide an efficient amplification in the energy of the response. We showcase this optimisation procedure on two flows: plane Poiseuille flow at Reynolds number $Re = 4000$, and turbulent flow past a NACA 0012 aerofoil at $Re = 23\,000$. In both cases, the optimisation yields sparse forcing modes that maintain important features of the structures arising from an SVD in order to provide a gain in energy. These results showcase the benefits of utilising a sparsity-promoting resolvent formulation to uncover sparse forcings, specifically with a view to using them as actuation locations for flow control.

Key words: shear-flow instability, instability control

1. Introduction

Resolvent analysis is a framework in which harmonic forcings that provide maximal amplification in their harmonic response can be determined on a frequency-by-frequency basis (Farrell & Ioannou 1993; Trefethen *et al.* 1993). By sweeping through frequencies, structural mechanisms that provide efficient means of flow amplification, as well as

† Email address for correspondence: c.s.skene@leeds.ac.uk

‡ Present address: Department of Applied Mathematics, University of Leeds, Leeds LS2 9JT, UK.

§ Present address: Department of Mechanical and Aerospace Engineering, North Carolina State University, NC 27695, USA.

¶ Present address: Physical Science and Engineering Division, King Abdullah University of Science and Technology, Thuwal 23955-6900, Kingdom of Saudi Arabia.

effective frequencies at which to provide such forcings, can be identified. While the original resolvent analysis focused on perturbation dynamics about steady states, recent studies have extended the analysis to system dynamics about the mean-flow, with emphasis on examining the self-sustaining fluctuations that are characteristic of turbulent flows (McKeon & Sharma 2010). With the resolvent analysis being able to reveal the input–output relationship with respect to the chosen base state (Jovanović & Bamieh 2005), it serves naturally as a valuable tool to design flow control techniques. Past studies, including Luhar, Sharma & McKeon (2014), Yeh & Taira (2019), Toedtli, Luhar & McKeon (2019) and Liu *et al.* (2021), have demonstrated that physical insights revealed from resolvent analysis provide valuable guidance on developing effective and efficient actuation strategies.

Traditionally, modal analysis techniques for fluid flows (Taira *et al.* 2017, 2020) have been founded on L_2 -based norms, which can lead to global spatial structures. For the resolvent analysis, this translates to having forcing modes that are supported spatially over a large region. It should, however, be realised that actuation inputs cannot be applied over a large spatial region in practical engineering flow control settings. In general, flow control inputs can be introduced only as localised actuation inputs. To address this point, we consider sparsity-promoting approaches to target specifically resolvent forcing modes that have spatially compact support, i.e. are spatially sparse. We also note that sparsity-promoting techniques may also help to identify appropriate variables through which control inputs can be added to the flow for enhanced amplification. This piece of information is important in selecting the appropriate type of actuators to introduce control input to the flow field (Cattafesta & Sheplak 2011).

To sparsify the resolvent forcing mode, we adopt an approach similar to that of Foures, Caulfield & Schmid (2013), who used alternative norms for studying transient growth in plane Poiseuille flow. In their work, transient growth analysis has been treated as a gradient-based optimisation problem, where the goal is to find the initial condition that has the most growth as measured by an appropriate norm. Choosing the L_2 -norm leads to the usual transient growth analysis (Trefethen *et al.* 1993) that can be solved using a singular value decomposition. However, choosing an alternative norm can tune the analysis to reveal different mechanisms that would be sub-optimal in terms of the L_2 -norm.

Foures *et al.* (2013) found more localised transient growth mechanisms using the infinity-norm, i.e. by measuring the norm of the response by its maximum value rather than energy. The result of this is that the identified initial conditions are localised spatially in order to achieve responses that are focused around local ‘hot spots’. Further to this, the non-convex nature of this optimisation problem means that there exist different branches of optimal initial conditions, with some representing local maximums of the cost functional. Physically, these localisations manifested themselves in the form of initial conditions that focused either in the middle of the channel or towards the walls.

Following this approach, our study considers resolvent analysis as an optimisation problem where forcing modes are sought that maximise a prescribed cost functional. In order to obtain spatially sparse forcing modes, we propose a gradient-based algorithm that maximises the energy of the output whilst minimising the L_1 -norm of the forcing, which is also constrained to have unit energy. To provide an initial assessment of our proposed method, we consider two examples. First, we consider the same plane Poiseuille set-up as in Foures *et al.* (2013), allowing us to assess qualitatively the differences between localisation strategies for initial conditions and for forced problems. Second, we consider turbulent flow past an aerofoil using the same mean-flow as Yeh & Taira (2019), providing an assessment of the method in a higher-Reynolds-number turbulent flow.

The structure of the paper is as follows. Section 2 outlines the mathematical formulation of the paper and contains an introduction to the resolvent operator, background on Riemannian optimisation and how we utilise it to find optimal sparse resolvent modes, and a discussion of wavemakers in the context of a resolvent analysis. In § 3, we discuss the numerical set-up, with the results being presented subsequently in § 4. Conclusions are offered in § 5.

2. Formulation

2.1. The resolvent operator

Let us consider the Navier–Stokes equations in the general, spatially discretised form

$$\mathbf{G} \frac{dq}{dt} = \mathcal{N}(q), \tag{2.1}$$

where q is the spatially discretised state vector, and \mathcal{N} represents the right-hand side of the unforced Navier–Stokes equations. Including the mass-matrix \mathbf{G} in (2.1) means that this form could represent either the compressible Navier–Stokes equations or the incompressible equations where there is no time derivative in the continuity equation. By linearising this equation about a base-flow q_0 , we can write the system in input–output form as

$$\mathbf{G} \frac{dq}{dt} = \mathbf{L}_{q_0} q + \mathbf{B} f, \tag{2.2}$$

$$y = \mathbf{C} q, \tag{2.3}$$

where \mathbf{L}_{q_0} is the linearised Navier–Stokes operator (Jeun, Nichols & Jovanović 2016). The matrix \mathbf{B} allows for the introduced forcing f (input) to be windowed in space or restricted to specific equations or state variables. In an analogous manner, the matrix \mathbf{C} allows for a similar windowing to be applied to the output y .

The relationship between harmonic inputs and outputs with frequency ω can be obtained by Laplace transforming the input–output system in time, giving the relation

$$\hat{y} = \mathbf{C}(-i\omega \mathbf{G} - \mathbf{L}_{q_0})^{-1} \mathbf{B} \hat{f}. \tag{2.4}$$

Through this equation, the resolvent operator is defined via $\mathcal{H}_{q_0} \equiv \mathbf{C}(-i\omega \mathbf{G} - \mathbf{L}_{q_0})^{-1} \mathbf{B}$. The form of (2.4) shows that the resolvent operator is equivalent to a transfer function that maps the forcing to its time-asymptotic response. Before we discuss the meaning of the resolvent in fluid dynamics, it is worth considering the Laplace variable ω . If the operator \mathbf{L}_{q_0} is stable, then ω is real and (2.4) is obtainable via the Fourier transform. However, if \mathbf{L}_{q_0} is unstable, then more care is needed. Indeed, for unstable \mathbf{L}_{q_0} , the time-asymptotic response is not given via (2.4) and is instead a combination of the exponentially growing disturbance given by the most unstable eigenvector and the forced response given by the resolvent. In order to separate these two mechanisms, a complex value for ω can be used, leading to the concept of a time-discounted resolvent analysis (Jovanović 2004). Choosing complex values for ω means that the imaginary part can be chosen such that the forced response ‘rises above’ the exponentially growing disturbance due to the unstable nature of \mathbf{L}_{q_0} , allowing for the forced dynamics to be probed (see Yeh *et al.* (2020) for more details).

In the context of fluid dynamics, the resolvent can be interpreted in two main ways. First, choosing q_0 to be a steady solution of the unforced Navier–Stokes equations leads to a non-normal stability study of the flow. In this manner, the resolvent identifies forcing

structures that lead to particularly efficient amplification in the dynamics despite the stable nature of the flow (Trefethen *et al.* 1993). Second, using a time-averaged mean-flow for q_0 leads to the resolvent formulation of turbulence (McKeon & Sharma 2010). The resolvent in this instance can be used to identify the coherent structures that arise via disturbances caused by the nonlinear terms.

For both steady base-flows and time-averaged mean-flows, the resolvent provides critical insights into how forcings can cause an amplification in the dynamics. This amplification can occur both from resonant frequencies and also from particularly effective structural mechanisms. Whilst one could choose a variety of forcings \hat{f} at each frequency to determine the most effective structures, it is more efficient to solve directly for the optimal forcing. This can be formulated mathematically as

$$f_{opt} = \arg \max_f \frac{\|y\|_{w_q}}{\|f\|_{w_f}}, \tag{2.5}$$

where the norms are defined as $\|a\|_W^2 = a^H W a$, with W being a positive definite weight matrix. We allow for the weight matrices for the forcing (W_f) and response (W_q) to be different. These matrices are problem-dependent, and are chosen so that the norms represent appropriate measures of the energy (see §§ 3.1 and 3.2 for examples). The cost functional in this case is the gain. To link the weighted norms to the two norms, it is useful to also consider the Cholesky decomposition $W = M^H M$. The optimal forcing has the corresponding output $\tilde{y}_{opt} = \mathcal{H} f_{opt}$, with the amount of amplification being measured by the gain $\sigma = \|\tilde{y}_{opt}\|_{w_q} / \|f_{opt}\|_{w_f}$. This problem can be solved by taking the singular value decomposition (SVD) of $M_q \mathcal{H} M_f^{-1}$, whose maximum singular triplet is $(\sigma, M_q y_{opt}, M_f f_{opt})$.

While a resolvent analysis in this manner can provide useful information about frequencies and forcing structures that can provide a large amplification, and therefore identify good candidates for flow control, the forcing structures are often global. This means that implementing them in a practical situation is infeasible. In the present study, we present an approach to finding sparse (spatially compact) resolvent forcings that induce large amplifications in the underlying dynamics. In this manner, particularly sensitive spatial locations in the flow field are identified, providing a guide for effective and efficient actuation.

2.2. Sparsification via Riemannian optimisation

To seek a spatially sparse resolvent forcing mode, we first generalise the optimal forcing problem. We start by realising that finding the greatest singular value of the resolvent matrix is equivalent to maximising the gain

$$\sigma = \frac{\|\mathcal{H}f\|_{w_q}^2}{\|f\|_{w_f}^2}. \tag{2.6}$$

Therefore, instead of carrying out an SVD, we could instead maximise the gain via a gradient ascent algorithm. It is useful to phrase this optimisation as follows: maximise the gain

$$\sigma = \|\mathcal{H}f\|_{w_q}^2, \tag{2.7}$$

where the forcing is confined to the manifold given by the constraint $\|f\|_{w_f}^2 = 1$. This problem is equivalent to (2.6) because the resolvent is linear and hence will produce the

same gain defined by (2.6) if we choose the forcing to have a unit-energy norm. In effect, by constraining our forcing to this manifold, we are ensuring that we search for the maximum amplification in dynamics with the forcing having the same energy budget.

Whilst we could conduct an unconstrained optimisation by enforcing $\|f\|_{W_f}^2 = 1$ with a Lagrange multiplier (Pringle, Willis & Kerswell 2012), we instead take account of this constraint directly in the update step. This results in an approach similar to that of Foures *et al.* (2013), where a geometric approach (Douglas, Amari & Kung 1998) was used to ensure that the unit-norm condition is satisfied when stepping in the search direction. In general, carrying out an optimisation where the input is constrained to a manifold is known as Riemannian optimisation (Absil, Mahony & Sepulchre 2007).

Let us first discuss the optimisation problem considered thus far. We seek to maximise the gain

$$\sigma^{(L_2)} = \|M_q \mathcal{H} M_f^{-1} M_f f\|_2, \tag{2.8}$$

subject to $\|M_f f\|_2 = 1$. By expressing the problem in this form, we have reformulated the problem in terms of the L_2 -norm, hence we are optimising with respect to the vector $f_M = M_f f$, which we constrain to have unit L_2 -norm. The manifold for this problem then becomes the complex hypersphere $\mathcal{S} = \{y \mid y^H y = 1\}$.

For an unconstrained optimisation, we generally work with the Euclidean gradient

$$\nabla_{f_M} \sigma^{(L_2)} = \frac{\partial \sigma^{(L_2)}}{\partial f_M}. \tag{2.9}$$

By stepping in the direction of the conjugate of this gradient, we would be increasing the value of $\sigma^{(L_2)}$, assuming that we use a sufficiently small step size for which a linear approximation is appropriate. The problem with this approach is that stepping in such a direction would most likely result in a vector that is no longer on the manifold.

To carry out a gradient descent on the hypersphere, we must therefore define the gradients appropriately. Riemannian optimisation will not work directly with the Euclidean gradient, but instead all gradients must be tangent to the hypersphere at the evaluation points. The set of all vectors tangent to the manifold at a point x is known as the tangent space $T_x \mathcal{S}$, with the set of all tangent spaces being referred to as the tangent bundle $T\mathcal{S} = \sum_{x \in \mathcal{S}} T_x \mathcal{S}$ (see figure 1, which shows schematically the Riemannian optimisation procedure). For the hypersphere, the Riemannian gradient can be written as

$$\text{grad } \sigma^{(L_2)}(f_M) = \nabla_{f_M} \sigma^{(L_2)} - \left(f_M^H \nabla_{f_M} \sigma^{(L_2)} \right) f_M = \text{Proj}_{f_M} \left(\nabla_{f_M} \sigma^{(L_2)} \right), \tag{2.10}$$

where the function Proj is the projection that links the Riemannian gradient to the Euclidean one.

Now that we have defined appropriate gradients, we must also define how to step in the direction of steepest ascent. For the unconstrained optimisation, we may simply add a scalar multiple (the step size) of this gradient onto our current value of the forcing. However, for the Riemannian optimisation, this will result in a vector that no longer falls on the manifold, as noted above. The equivalent procedure in this case is the notion of a retraction. A retraction is a mapping $R_x(\xi) : T_x \mathcal{S} \rightarrow \mathcal{S}$ such that $R_x(\mathbf{0}) = x$ and $DR_x(\mathbf{0}) = \text{id}_{T_x \mathcal{S}}$. In other words, a retraction maps vectors tangent to the manifold at x to other vectors on the manifold such that for $\xi = \mathbf{0}$ it maps x to itself, and such that the derivative of the mapping at $\xi = \mathbf{0}$ is the identity. For the hypersphere, we have the




Figure 1. A sketch illustrating the concept of Riemannian optimisation. First, the Euclidean gradient $\nabla_{f_M^n} \sigma^{(L_2)}$ is found from the vector f_M^n that is situated on the hypersphere \mathcal{S} . This vector is then mapped to the tangent space $T_{f_M^n} \mathcal{S}$ via the projection $\text{Proj}_{f_M^n}$. Next, the Riemannian gradient is extended along the tangent space by the step size α . Finally, we map this gradient back to the manifold via the retraction $R_{f_M^n}$, yielding the updated forcing vector f_M^{n+1} . For varying values of α , the retraction traces out a smooth curve over the manifold, displayed as a dotted line. The link between α and θ is also shown.

retraction

$$R_x(\xi) = \frac{x + \xi}{\|x + \xi\|}. \tag{2.11}$$

Once the gradient is found, we can then update the forcing using the map $R_{f_M}(\alpha \text{grad } \sigma^{(L_2)}(f_M^n))$, where α denotes the step size. By writing $\cos(\theta) = 1/\sqrt{1 + \alpha^2}$, we can also express the update step as

$$f_M^{n+1} = R_{f_M}(\alpha \text{grad } \sigma^{(L_2)}(f_M^n)) = \cos(\theta) f_M^n + \sin(\theta) \text{grad } \sigma^{(L_2)}(f_M^n), \tag{2.12}$$

which is exactly the geometric form used by Foures *et al.* (2013). Note that we have described a steepest ascent approach here. However, many alternative gradient-based optimisation algorithms, such as the conjugate gradient and Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithms, are applicable to Riemannian optimisation with faster convergence (Boumal & Absil 2015; Huang, Absil & Gallivan 2016).

The main advantage of phrasing the optimal forcing output problem in this way is its generality. Whilst we have shown how we can obtain the same result as the SVD (and it is actually possible to get the higher-order singular values in this manner by considering a different manifold), we are free to change how we define the gain. The SVD can find the gain only in the L_2 -norm sense. This means that the input is measured by an energy norm, leading to the global structures seen in many studies. In order to introduce sparsification, we consider the use of the 1-norm.

A sketch of the unit L_2 - and L_1 -norms for a vector (x, y) in \mathbb{R}^2 is presented in figure 2(a). The L_2 -norm takes the form of a circle, whereas the L_1 -norm yields a regular diamond inscribed within this circle. Note that the unit L_1 -norm touches the unit L_2 -norm at the coordinate axes. This indicates that the L_1 -norm for all vectors with unit L_2 -norm yields its smallest value for sparse vectors, i.e. vectors (x, y) with either x or y equal to zero. Indeed, if the diamond touched the circle at another location (x_0, y_0) with $x_0 \neq 0$ and $y_0 \neq 0$, then the L_2 -norm would be unity whereas the L_1 -norm would have the value $|x_0| + |y_0| > 1$. Hence optimising over the space of unit-norm forcings, whilst penalising the L_1 -norm, will push the forcing vector, and hence its structure, towards more locally supported structures.

Sparsifying the resolvent forcing mode




Figure 2. Sketches of the L_2 - and L_1 -norms. The effect of a coordinate rotation on the norms is also demonstrated. (a) The unit L_2 -norm (circle) and L_1 -norm (diamond) with respect to the coordinate system (x, y) . (b) The L_2 -norm (green dashed circle) and L_1 -norm (green dashed square) with respect to the coordinate system (ξ, η) .

One important consideration when using an alternative norm, such as the L_1 -norm, is illustrated by figure 2(b). Here, we have again shown the unit L_2 - and L_1 -norms, but this time for the coordinate system (η, ξ) that is obtained via a rotation of the coordinate system (x, y) by an angle θ . In this new coordinate system, the L_2 -norm still represents a circle, which is invariant under this transformation. However, the unit L_1 -norm is affected, and its square shape is rotated by the angle θ . This means that in this new coordinate system, the sparse vectors, where the square touches the circle, are different to those of the original coordinate system (x, y) . In other words, what is considered sparse is defined completely by how we choose to represent our vectors. In practice, we must be careful when choosing the vector of which we take the L_1 -norm. We therefore choose to take the L_1 -norm of a vector that leaves the L_2 -norm unchanged, yet has appropriate axes for best defining the sparsity that we aim to achieve. In terms of resolvent analysis, this transformation is used to maintain the physical relevance of the sparsification. Specific examples are described in §§ 3.1 and 3.2.

Based on the discussion of the previous two paragraphs, we seek to maximise the new gain $\sigma^{(L_1)}$ defined by

$$\sigma^{(L_1)} = \frac{\sigma^{(L_2)}}{\|T(\mathbf{f}_M)\|_1} = \frac{\|M_q \mathcal{H} M_f^{-1} \mathbf{f}_M\|_2}{\|T(\mathbf{f}_M)\|_1}, \quad (2.13)$$

still subject to the forcing \mathbf{f}_M having a unit-energy norm. The transformation T in the denominator is a transformation of the vector \mathbf{f}_M to another vector. Hence the vector in the denominator need not be equal to the forcing vector \mathbf{f}_M as, based on the discussion of the previous paragraph, this may not be relevant physically. However, by ensuring $\|T(\mathbf{f}_M)\|_2 = \|\mathbf{f}_M\|_2$, we maintain the geometric interpretation of sparsity illustrated by figure 2(a), albeit in a much higher-dimensional space. By dividing the usual gain by the 1-norm of the vector $T(\mathbf{f}_M)$, we are in effect promoting sparsity, with sparsity defined as a vector $T(\mathbf{f}_M)$ with a minimal number of non-zero entries. Optimising the gain in this form will seek a compromise between providing a large gain in energy whilst ensuring the spatial sparsity of the forcing. Indeed, the maximal nature of $\sigma^{(L_1)}$ means that obtaining a response with larger energy requires a forcing structure that is less sparse. Likewise, making the forcing more sparse leads to a less energetic response.

As in the study of Foures *et al.* (2013), who considered a similar optimisation problem for localising flow structures obtained in transient growth studies, our cost functional is

non-convex. This means that any solution to the optimisation problem is guaranteed only to be a local, rather than global, maximum of the cost functional. In the case of transient growth, this led to multiple branches of solutions being found during the optimisation, which could be discovered by running the problem with multiple starting guesses for the gradient-based optimisation. However, despite running multiple instances of each optimisation with different initial guesses in our examples, no differences in the solution could be found apart from symmetries of the flow, which are to be expected. Whilst this does not confirm that our results are truly the global optimum, it does highlight a difference between localising forcings for driven versus initial-condition based studies.

Another important factor is the realisation that the L_1 -norm is notoriously hard to optimise due to its non-smoothness near the origin. Intuitively, we can visualise the problem by considering the unconstrained optimisation problem of minimising the L_1 -norm of a scalar a . Using our gradient-based approach, this amounts to stepping in the direction of steepest descent, which for our simple example is the sign of a . No matter how near or far we are to the optimal value of $a = 0$, this gradient will have the same magnitude. This means that we will step continuously over the optimal value, unless the step size is perfect, leading to zig-zagging and ultimately causing the algorithm to converge rather slowly. To alleviate this behaviour we replace the L_1 -norm with a smooth approximation, namely $l_{1,\delta}(\mathbf{q}) = h_\delta(\mathbf{q})/\delta$, where $h_\delta(\mathbf{q})$ is the pseudo-Huber-norm (Bube & Langan 1997; Bube & Nemeth 2007)

$$h_\delta(\mathbf{q}) = \sum_j \delta^2 \left(\sqrt{1 + \frac{|q_j|^2}{\delta^2}} - 1 \right). \tag{2.14}$$

This pseudo-Huber-norm has the property that it approximates the L_1 -norm for small δ and is completely smooth. Therefore, in order to achieve convergence, we will not optimise $\sigma^{(L_1)}$ directly but perform a series of optimisations for the quantity

$$\sigma_\delta^{(L_1)} = \frac{\|\mathbf{M}_q \mathcal{H} \mathbf{M}_f^{-1} \mathbf{f}_M\|_2}{l_{1,\delta}(\mathbf{T}(\mathbf{f}_M))}, \tag{2.15}$$

for decreasing values of δ . By using the optimal forcing obtained from an optimisation for the preceding one with a lower value of δ , we are able to achieve more robustly a converged optimisation for a sufficiently small δ such that our norm (2.14) is an appropriate approximation for the true L_1 -norm.

Before concluding this section, it is important to note that our choice of cost functional is not unique. Indeed, other cost functionals such as $\sigma^{(L_1)} = \sigma^{(L_2)} - \mu \|\mathbf{T}(\mathbf{f}_M)\|_1$ can also lead to sparse forcing modes for appropriate choices of μ . However, the fact that unit-norm forcings can lead to gains in energy many orders of magnitude larger than that of the forcing makes the choice of μ , which must balance the L_2 -based gain against the L_1 -based forcing, a difficult challenge. This is complicated further by the strong dependence of the gain on the forcing frequency, making a universally good way of choosing μ hard to determine. In our proposed cost functional there is no such parameter to choose, meaning that it can be applied easily to different frequencies and base-flows without change. Hence we continue with it for the rest of the study.

2.3. Resolvent wavemaker

One concept that we use in our subsequent analysis is that of structural sensitivity and the wavemaker (Giannetti & Luchini 2007). The wavemaker has its origin in global stability

analysis and provides a way to highlight regions in which flow field changes result in changes to global modes. Specifically, the wavemaker is the structural sensitivity to a localised flow feedback. To obtain the wavemaker, we consider the eigenvalue problem $\mathbf{L}\mathbf{x} = \lambda\mathbf{G}\mathbf{x}$. This eigenvalue problem could arise, for instance, as a global stability problem, in which case \mathbf{x} would be the global mode, with the corresponding eigenvalue λ giving its frequency and growth rate. It can be shown (see the review of Schmid & Brandt (2014), for example) that to first order, a perturbation to the eigenvalue $\delta\lambda$ for a perturbation in the matrix $\delta\mathbf{L}$ is given via

$$\delta\lambda = \frac{\langle \mathbf{x}^\dagger, \delta\mathbf{L}\mathbf{x} \rangle}{\langle \mathbf{x}^\dagger, \mathbf{G}\mathbf{x} \rangle}, \quad (2.16)$$

where \mathbf{x}^\dagger is the solution to the adjoint eigenvalue problem $\mathbf{L}^H\mathbf{x}^\dagger = \bar{\lambda}\mathbf{G}^H\mathbf{x}^\dagger$. The wavemaker is then obtained by specifying $\delta\mathbf{L} = \mathbf{I}$ and instead taking the elementwise, or Hadamard (\odot), product:

$$\lambda = \frac{\bar{\mathbf{x}}^\dagger \odot \mathbf{x}}{\langle \mathbf{x}^\dagger, \mathbf{G}\mathbf{x} \rangle}. \quad (2.17)$$

In this way, the wavemaker λ can then be thought of as a vector field $\lambda(x, y) = (\lambda_u, \lambda_v, \lambda_w)$ whose components represent what changes to the eigenvalue occur from localised feedback at each location and state component in the flow field.

We note quickly that there are a few ways in which the wavemaker could be perceived. Whilst we have stayed within a discrete setting, Giannetti & Luchini (2007) present the wavemaker in a continuous formulation. This gives the main difference that their wavemaker is a scalar field that is defined pointwise via $\lambda(x, y) = \|\mathbf{x}^\dagger(x, y)\| \|\mathbf{x}(x, y)\|$. Hence their wavemaker, by the Cauchy–Schwarz inequality, shows the maximum change to the eigenvalue that can be achieved via localised feedback at each spatial location. Conversely, the wavemaker presented by Schmid & Brandt (2014) is more easily related to ours via $\lambda(x, y) = \lambda_u + \lambda_v + \lambda_w$. In this manner, they obtain a complex-valued wavemaker whose real and imaginary parts show the individual changes to the real and imaginary parts of the eigenvalue. Additionally, the sign of these changes is retained, allowing for the direction of the eigenvalue perturbation to be determined. However, by keeping the values of the flow field separate, our wavemaker definition is strongly related to that of Paladin *et al.* (2019), who introduce windowing matrices to allow for the selection of specific physical components in the resulting wavemaker. Whilst they use these matrices to isolate the contribution of the momentum to the wavemaker, we instead do this procedure for each separate velocity component. This means that for each spatial location, our wavemaker tells us how a specific eigenvalue will move for localised feedback restricted to each component of the state vector.

Whilst the previous paragraph talked about a wavemaker in terms of an eigenvalue problem, it can also be directly formulated for an SVD-based resolvent analysis (Qadri & Schmid 2017). Indeed, by realising that taking the SVD of the matrix $\mathbf{K} = \mathbf{M}_q \mathcal{H} \mathbf{M}_f^{-1}$ is equivalent to taking the eigenvalues of the matrix $\mathbf{K}^H \mathbf{K}$, the same procedure that yields (2.16) can be applied, resulting in

$$\delta\sigma = \sigma^2 \text{Re}(\langle \mathbf{f}, \delta\mathbf{L}\mathbf{q} \rangle \mathbf{w}_f), \quad (2.18)$$

where \mathbf{L} stands for the linearised Navier–Stokes operator (Fosas de Pando, Schmid & Lele 2014; Fosas de Pando & Schmid 2017; Qadri & Schmid 2017). Again, taking $\delta\mathbf{L} = \mathbf{I}$ and




Figure 3. The eigenvalue-based wavenumbers and resolvent wavenumbers for cylinder flow at $Re = 100$ shown for illustration. Computations are performed with the immersed boundary projection method *ibmos* (Fosas de Pando 2020). (a) Wavenumber for the real part of the u -component. (b) Wavenumber for the real part of the v -component. (c) Resolvent wavenumber for the u -component at $St = 0.162$. (d) Resolvent wavenumber for the v -component at $St = 0.162$.

using the Hadamard product yields

$$\sigma = \sigma^2 \text{Re}(\tilde{f} \odot \mathbf{W}_f \mathbf{q}). \tag{2.19}$$

The resolvent wavenumber σ is then analogous to the eigenvalue-based wavenumber, i.e. for localised feedback at each spatial location and component of the state vector, the resolvent wavenumber will indicate how the singular value will be perturbed.

An example of the wavenumber and resolvent wavenumber is shown for cylinder flow in figure 3. It is important to note that for the eigenvalue-based wavenumber, the frequency is set by the eigenvalues. However, our definition of the resolvent wavenumber allows any frequency to be specified. Therefore, we concentrate on $St = 0.162$, which is the frequency at which the most unstable eigenvalue is found. We observe that these wavenumbers have similar structures but with different gains. The fact that they have similar structures is not surprising, since the resolvent forcing and response modes are similar qualitatively to the direct and adjoint eigenvectors, respectively. However, the signs of the structures are often different. This indicates that a localised feedback affects the eigenvalue perturbation differently from the singular value perturbation, highlighting the importance of formulating a resolvent wavenumber in order to quantify the effect of localised feedback for resolvent analyses.

3. Numerical set-up

This section describes the numerical set-up for our flow examples. In addition to the details given in this section, all Riemannian optimisations are managed using the Python package *pyManopt* (Townsend, Koep & Weichwald 2016), the Python extension of the MATLAB

package *Manopt* (Boumal *et al.* 2014). The optimisations are all conducted using the conjugate gradient algorithm.

3.1. Plane Poiseuille flow

First, we consider plane Poiseuille flow to compare the differences in localisation strategies between initial-condition-based (transient growth) and driven (resolvent) studies. The present set-up follows that of Foures *et al.* (2013) in which a domain covers $(x, y) \in [0, 2\pi] \times [0, 2]$ at Reynolds number $Re = 4000$. No-slip boundary conditions are applied at $y = 0$ and 2 , and periodic boundary conditions are applied at $x = 0$ and 2π . The base-flow is provided analytically as $u = y(2 - y)$. We conduct the numerical simulations using the Python package *ibmos* developed by Fosas de Pando (2020). This is an immersed boundary projection code based on the formulation of Taira & Colonius (2007) with specific formulation for optimisation and stability analyses. The package solves the nonlinear incompressible Navier–Stokes equations and provides directly the linearised and adjoint codes necessary for conducting a resolvent analysis. For the plane Poiseuille examples, the matrix \mathbf{B} is chosen so that the forcing is added to only the momentum equations. Similarly, the matrix \mathbf{C} is chosen so that only velocity components constitute the output.

As detailed in § 2.2, there is some consideration in choosing the vector for our L_1 -norm, $T(\mathbf{f}_M)$. The obvious choice would be to use the same vector used for the unit-energy norm, \mathbf{f}_M , for the L_1 -norm. However, as the x - and y -components of the velocity occur in different locations of \mathbf{f}_M , this would result in a sparsification that not only sparsifies the forcing mode in space, but also sparsifies between the x - and y -components of velocity. In other words, if the sparse procedure were to locate a single spatial point for the forcing mode, then it would also be advantageous to align the velocity vector completely with the coordinate axes at this point, in order to achieve a further reduction in the L_1 -norm. As we are interested primarily in localisation in space, as opposed to sparsifying the velocity vector itself, we therefore design a vector for the L_1 -norm optimisation that does not result in this unwanted sparsification. This is particularly pertinent to applications of this method to flow-actuation, where the directional information obtained by keeping the x -, y - and possibly z -components of velocity independent of the sparsification procedure will provide additional insight into actuator design.

To this end, we consider a vector of the form $T(\mathbf{f}_M) = \mathbf{M}(\mathbf{u} \odot \mathbf{u} + \mathbf{v} \odot \mathbf{v})^{1/2}$, where \odot is the Hadamard product, and the square root is taken componentwise. This vector has the same 2-norm as \mathbf{f}_M , but groups local contributions of the forcing mode to the total energy together. Hence the L_1 -norm of this vector is small when the forcing is localised in space, but without penalising among individual components of the velocity vector. It should be noted that some additional care may be needed when designing this vector, depending on the specific numerical implementation. For example, as our immersed boundary implementation uses a staggered mesh with the x -components of velocity lying on the east and west faces of the cell whilst the y -components lie on the north and south faces, we form the vector $T(\mathbf{f}_M)$ on the cell centres by averaging the kinetic energy contributions from the cell edges. The weight matrices are chosen to incorporate the grid spacing (see Taira & Colonius (2007) for more information), so that the forcing and response are measured in terms of the kinetic energy.

3.2. Flow past an aerofoil

Next, we also consider a spanwise-periodic turbulent flow over a canonical aerofoil obtained from a large-eddy simulation (LES) with a Vremen subgrid scale model

(Vreman 2004). The LES is conducted using the finite-volume solver *CharLES* that solves the compressible Navier–Stokes equations with second-order spatial and third-order temporal accuracies (Khalighi *et al.* 2011; Brès *et al.* 2017). The linearisation is performed within the same solver (Sun *et al.* 2017), considering the time- and spanwise-averaged turbulent flow over the aerofoil as the base-flow.

The resolvent analysis is performed on a mesh separate from that used by the LES. The mesh for the resolvent analysis has a two-dimensional rectangular domain of extent $x/L_c \in [-15, 16]$ and $y/L_c \in [-6, 5]$, comprising approximately 0.11 million cells and giving the resulting discretised linear operator dimension $540\,840 \times 540\,840$. Compared to the LES mesh, the mesh for resolvent analysis is coarser over the aerofoil and in the wake, but is much finer in the upstream region of the aerofoil in order to resolve the forcing mode structures. The convergence of resolvent norm with respect to the domain extent and grid resolution has been reported in detail in Yeh & Taira (2019). At the far-field boundary and over the aerofoil, Dirichlet conditions are specified for the density and velocities, and Neumann conditions are prescribed for the pressure in \mathbf{q} . At the outlet boundary, Neumann conditions are provided for all flow variables. The base-flow is two-dimensional, however, in contrast to the plane Poiseuille case, and we allow the perturbations to be three-dimensional by adopting a bi-global setting that decomposes \mathbf{q} into spanwise Fourier modes with wavenumber β .

Even though the linear operator is sparse, its large dimension requires special care. To deal efficiently with this operator, the Python bindings for PETSc (Balay *et al.* 1997, 2021a, b), petsc4py (Dalcin *et al.* 2011) are used. This enables us to carry out the required linear algebra manipulations in parallel whilst keeping our code within the Python environment. Specifically, PETSc is used together with the external library MUMPS (Amestoy *et al.* 2001, 2019) in order to provide the LU decomposition of the resolvent operator, and hence evaluate the action of the resolvent operator (and its adjoint) on a vector. To compare our sparse method with a traditional resolvent analysis, the SVD of the resolvent operator is found using a Lanczos SVD solver provided by the Python bindings for SLEPc (Hernandez, Roman & Vidal 2005), slepc4py (Dalcin *et al.* 2011). The numerical implementation for taking the SVD of the resolvent has been made available (Skene, Ribeiro & Taira 2022).

As in the previous case, we need to be careful regarding the choice of the state vector for the L_1 -norm. To sparsify the location of any momentum input, rather than the individual components of the momentum, we must design our L_1 -norm such that the momentum components are together. This requires some care for the aerofoil case, since it is compressible and the modes are measured not via an L_2 -norm but via the Chu-norm (Chu 1965)

$$\|\mathbf{q}\|_E^2 = \int_{\Omega} \left(\frac{RT_0}{\rho_0} |\rho'|^2 + \rho_0 \|\mathbf{u}'\|^2 + \frac{R\rho_0}{(\gamma - 1)T_0} |T'|^2 \right) dV, \quad (3.1)$$

which represents the energy contained in a perturbation in the absence of compression work. In defining the Chu-norm, we have used a dash ' to denote quantities derived from our state vector \mathbf{q} . Similarly, a subscript 0 is used to denote quantities derived from the base-flow used for linearisation. This integral is discretised to $\|\mathbf{q}\|_E^2 = \mathbf{q}^H \mathbf{W}_E \mathbf{q}$, with \mathbf{W}_E as a positive definite weight matrix. Taking the Cholesky decomposition $\mathbf{W}_E = \mathbf{M}^H \mathbf{M}$ gives the matrices needed for the resolvent description ($\mathbf{M}_q = \mathbf{M}_f = \mathbf{M}$). Hence, to keep momentum grouped in our sparsification, we split the components of the norm matrix \mathbf{M} to form the state $T(\mathbf{f}_M) = (\mathbf{M}_\rho \rho', \mathbf{M}_{KE} \kappa', \mathbf{M}_T T')$, where $\kappa' = \sqrt{(\rho \mathbf{u}') \odot (\rho \mathbf{u}') + (\rho v') \odot (\rho v') + (\rho w') \odot (\rho w')}$. Note that in defining κ' , we have used

the notation $(\rho\mathbf{u})' = \rho_0\mathbf{u}' + \rho'\mathbf{u}_0$ for the streamwise linearised momentum component, with similar definitions for the spanwise $(\rho v)'$ and transverse $(\rho\mathbf{w})'$ linearised momentum components. This vector has the same L_2 -norm as our full state vector. However, when taking the L_1 -norm, the kinetic energy is now grouped, ensuring that all velocity components are treated equally. It should be pointed out that now we have two additional components in the norm, specifically ρ' and T' . By not including these in the same component of the vector for the L_1 -norm, they are treated separately by the sparsification. In essence, this means that the sparse resolvent not only sparsifies the spatial structure of any forcing but also sparsifies the actuation mechanism by choosing between a velocity-based, density-based or temperature-based forcing.

4. Results

4.1. Plane Poiseuille flow

Let us first consider plane Poiseuille flow. This canonical example provides a good comparison with the work of Foures *et al.* (2013) and highlights the differences between using an alternative norm for a resolvent analysis and a transient growth study. It is important to note that as plane Poiseuille flow is a parallel flow, we could have proceeded with a local analysis, i.e. we could specify the streamwise wavenumber α and search for modes of the form

$$\mathbf{f}'(x, y) = \mathbf{f}'_\alpha(y) e^{i\alpha x} \quad \text{and} \quad \mathbf{u}(x, y) = \mathbf{u}_\alpha(y) e^{i\alpha x}. \quad (4.1a,b)$$

However, as we are using a global (two-dimensional) analysis, the wavenumbers that our forcing and response can consist of are set by the aspect ratio of the domain. Taking $x \in [0, 2\pi]$ with periodic boundary conditions requires our wavenumbers to be integer, i.e. $\alpha \in \mathbb{N}$. Another artefact of using a two-dimensional code for a parallel-flow is that the results do not change if the forcing and response modes are translated along the x -axis.

The gains obtained from a full resolvent analysis (i.e. by using an SVD) are shown in figure 4. This figure shows a strong peak at $\omega = 0.278$, followed by another peak at $\omega = 1.14$. Examining the forcing and response modes at these two frequencies (shown in figures 5 and 6, respectively), we observe that the first peak is associated with $\alpha = 1$ structures, whereas the second peak corresponds to a higher wavenumber, $\alpha = 2$. This can be seen as a consequence of the flow being parallel and hints that by performing a two-dimensional analysis, the optimal response is obtained at the wavenumber that has the maximum response from the one-dimensional analysis. With this in mind, the qualitative shape of the gain distribution agrees with those obtained in Schmid & Henningson (2001) if the effects of perturbations in the spanwise direction, not considered in our analysis, are neglected. In both cases, the forcing mode consists of structures slanted against the shear, indicating that an Orr mechanism is responsible for the gain in dynamics. Another interesting observation can be made by examining the phase velocity $k = \omega/\alpha$. We see that the phase velocity of the second peak is twice that of the first peak. The fact that the second peak is a faster disturbance is also evident from the forcing mode being situated more centrally in the y -direction where the base-flow has a higher velocity.

We now turn our attention to the resolvent analysis results from the sparse optimisation procedure. The sparse forcing mode obtained for $\omega = 0.278$ is shown in figure 7. First, it is clear from this figure that the forcing mode is more sparse than the full resolvent analysis. Indeed, instead of a series of slanted structures angled against the shear, we now have thin stripes parallel to the walls. Interestingly, even though our vector $\mathbf{T}(\mathbf{f}_M)$ was chosen carefully not to sparsify the separate velocity components at a given spatial




Figure 4. The optimal gains over frequency for Poiseuille flow. Our analysis focuses on the optimal gain that occurs at $\omega = 0.278$, and on the second peak at $\omega = 1.14$ (both shown with a black +).




Figure 5. The real parts of the forcing and response modes obtained by SVD of the resolvent at $\omega = 0.278$. The forcing is unit-norm, whereas the response mode has norm equal to the gain. (a) Real part of the full forcing mode, u -velocity component. (b) Real part of the full forcing mode, v -velocity component. (c) Real part of the full response mode, u -velocity component. (d) Real part of the full response mode, v -velocity component.

location, the sparse forcing mode consists mainly of a u -component, indicating that this forcing is primarily in the direction of the wall.

A striking feature of the sparse forcing mode is that it has maintained its $\alpha = 1$ structure, i.e. it is still 2π -periodic in the x -direction. This is particularly enlightening since the strip structure is not as sparse as the forcing mode could be, which would consist of just one element of the kinetic energy vector being filled, i.e. a single spatial location forcing. Therefore, the fact that the sparse procedure has chosen a less sparse structure indicates that forcing with this spatial wavenumber is crucial in achieving a high gain at this frequency. The location of these stripes can be hypothesised to be intrinsically linked to the $\alpha = 1$ structure using the concept of critical layers. A critical layer occurs at y^* where the base-flow velocity $U(y^*)$ is equal to the phase velocity k of a disturbance, and is central in causing instability in plane Poiseuille flow. Using the phase velocity k for a disturbance at $\omega = 0.278$ and $\alpha = 1$ implies a critical layer at $y^* \approx 0.150$, which is close to the y -location of the stripes, which occur at $y = 0.155$. Hence the sparse forcing

Sparsifying the resolvent forcing mode




Figure 6. The real parts of the forcing and response modes obtained by SVD of the resolvent at $\omega = 1.14$. The forcing is unit-norm, whereas the response mode has norm equal to the gain. (a) Real part of the full forcing mode, u -velocity component. (b) Real part of the full forcing mode, v -velocity component. (c) Real part of the full response mode, u -velocity component. (d) Real part of the full response mode, v -velocity component.




Figure 7. The real parts of the forcing and response modes obtained by sparsification at $\omega = 0.278$. The forcing is unit-norm, whereas the response mode has norm equal to the gain. (a) Real part of the sparse forcing mode, u -velocity component. (b) Real part of the sparse forcing mode, v -velocity component. (c) Real part of the sparse response mode, u -velocity component. (d) Real part of the sparse response mode, v -velocity component.

mechanism can be summarised as forcing along (or just above) and parallel to the critical layer, with the v -velocity component, which does not contribute to this critical layer, being negligible.

Further evidence for the importance of $\alpha = 1$ forcing is shown in the response modes, which are also displayed in figure 7. The figure shows that the response modes stemming from the sparse forcing mode have the same structure as those from the full resolvent. As well as reinforcing that the $\alpha = 1$ forcing is critical for providing optimal amplification at this frequency, this observation also highlights the low-rank nature of the resolvent at this frequency. Even though the forcing shape is qualitatively different in the sparse case, the shape of the response is identical, disregarding arbitrary phase shifts, albeit with a lower magnitude. This agrees with previous observations, such as that of Rosenberg, Symon & McKeon (2019), where it is shown that when there is a large separation in singular values, the shape of the forcing is less critical in exciting the dominant response.




Figure 8. The real parts of the forcing and response modes obtained using sparsification at $\omega = 1.14$. The forcing is unit-norm, whereas the response mode has norm equal to the gain. The location of the sparse forcing mode is shown with \times . (a) Real part of the sparse forcing mode, u -velocity component. (b) Real part of the sparse forcing mode, v -velocity component. (c) Real part of the sparse response mode, u -velocity component. (d) Real part of the sparse response mode, v -velocity component.

The lower magnitude is to be expected since our sparse forcing mode sacrifices some amount of energy to achieve a more localised spatial structure.

To provide a comparison between the results at different wavenumbers, we also carry out the sparse optimisation procedure at $\omega = 1.14$. Figure 8 shows the results, which differ quite significantly from the case of $\omega = 0.278$. In this case, the sparsification procedure has resulted in a single spatial forcing in u , with a negligible v -component that can be disregarded safely. In fact, the structure of the v -component is an artefact of the optimisation procedure, which initially converged to a critical-layer mechanism similar to the previous case, before converging to a single spatial location. The reason for the different structure in this case can be attributed to the higher-rank nature of the resolvent at this frequency. For $\omega = 0.278$, $\sigma_1^{(L_2)}/\sigma_2^{(L_2)} \approx 31$, whereas for $\omega = 1.14$, $\sigma_1^{(L_2)}/\sigma_2^{(L_2)} \approx 2$. The effect is that even though an $\alpha = 2$ forcing is optimal, there is a less clear distinction between this forcing and the higher-order singular vectors. The result is that, unlike the previous case, there is less of a need for a specific α wavenumber to provide the optimal gain, and the sparsification procedure can take advantage of this to sparsify the forcing structure further. This is also evident in the response modes, which are quite different from the SVD results. Finally, it is worth noting that the asymmetry of the forcing mode, with the single spatial location being located above the centreline, is due to our optimisation procedure converging to a local maximum. The reflection of this point about the centreline would also achieve the same value of the cost functional, hence representing another local optimum. Similar behaviour has been reported in the work of Foures *et al.* (2013).

4.2. Flow past an aerofoil

Next, we consider flow past a NACA 0012 aerofoil at angle of attack 9° , chord-based Reynolds number $Re = 23\,000$, and free stream Mach number $M = 0.3$ (see § 3.2 for the numerical details). In contrast to the plane Poiseuille example, this flow is unsteady and turbulent. Therefore, the mean-flow is used for linearisation. This time-averaged base-flow is shown in figure 9. Similarly to the work of Yeh & Taira (2019) and Ribeiro, Yeh & Taira (2020), who considered a resolvent analysis with the same base-flow, we consider

Sparsifying the resolvent forcing mode




Figure 9. The streamwise velocity component of the base-flow.




Figure 10. The optimal gains against the Strouhal number for flow past an aerofoil. Our analysis will focus on the optimal gain, which occurs at (a) $St \approx 5.22$ for $\beta = 0$, and (b) $St \approx 5.90$ for $\beta = 20\pi$ (both highlighted with a black +).

the resolvent modes at spanwise wavenumbers $\beta = 0$ and 20π . As the linear operator is unstable, a discounting parameter $\alpha = 0.63$ is used (Jovanović 2004). The gain–frequency relationships are shown in figure 10. Similarly to the previous examples, we focus our subsequent analysis on the frequencies at which the peak gain is obtained.

Let us begin our analysis by examining briefly the modes obtained from a full resolvent analysis for our chosen parameters. For full details, see the paper by Yeh & Taira (2019). The spanwise linearised momentum component of the forcing mode and its corresponding response for our two spanwise wavenumber choices are showcased in figure 11. In both cases, the forcing is similar, consisting of slanted structures near the leading edge of the aerofoil on the suction side. The response modes are both located in the shear layer further downstream of the leading edge, but differ in their spatial structures. For $\beta = 0$, there is a larger spatial support, with the mode shape extending both vertically and horizontally about the shear layer, whereas for $\beta = 20\pi$, the response aligns much more tightly with the shear. This agrees with the findings of Yeh & Taira (2019), who state that for an increased forcing frequency or wavenumber, the shear layer is needed to support the resulting smaller-scale fluctuations.

Now that we have characterised the L_2 -norm SVD-based results, we turn our attention towards the sparse-optimisation-based modes. Figure 12 shows the sparse forcing and response modes. In both spanwise wavenumber cases, the optimisation procedure has identified a single spatial momentum-based structure for the forcing mode, with the density




Figure 11. The full resolvent modes for aerofoil flow at $St = 5.22$ for $\beta = 0$, and $St = 5.90$ for $\beta = 20\pi$. The linearised component of the streamwise momentum is shown. (a) Full forcing mode for $\beta = 0$. (b) Full response mode for $\beta = 0$. (c) Full forcing mode for $\beta = 20\pi$. (d) Full response mode for $\beta = 20\pi$.




Figure 12. The sparse resolvent modes for aerofoil flow at $St = 5.22$ for $\beta = 0$, and $St = 5.90$ for $\beta = 20\pi$. The linearised component of the streamwise momentum is shown. (a) Sparse forcing mode for $\beta = 0$. (b) Sparse response mode for $\beta = 0$. (c) Sparse forcing mode for $\beta = 20\pi$. (d) Sparse response mode for $\beta = 20\pi$.

and pressure contributions being negligible. This illustrates the effectiveness of the sparse procedure, which in this case was not only able to sparsify the spatial structure, but has also sparsified the physical makeup of the forcing, indicating that a momentum-based forcing provides the optimal sparse gain in dynamics. It is also worth highlighting that even though we have only a single-location forcing, the response mode is qualitatively the same as the full case.

To provide additional insight into the chosen spatial location, we now examine the resolvent wavemakers, which are shown in figure 13. For both values of β , the wavemakers display a large positive region in the mean-flow shear layer. This is not unexpected, since regions of shear translate to regions of high non-normality in the linearised Navier–Stokes operator, which underpins sensitive areas for forcing. In fact, in the full-SVD forcing modes, we see this directly, as the forcing modes in both cases are located primarily in this shear region. The sparse forcing locations are also situated in this region and are

Sparsifying the resolvent forcing mode




Figure 13. The resolvent wavemakers for aerofoil flow at $St = 5.22$ for $\beta = 0$, and $St = 5.90$ for $\beta = 20\pi$. The linearised component of the streamwise momentum is shown. (a) Wavemaker for $\beta = 0$. (b) Wavemaker for $\beta = 20\pi$.

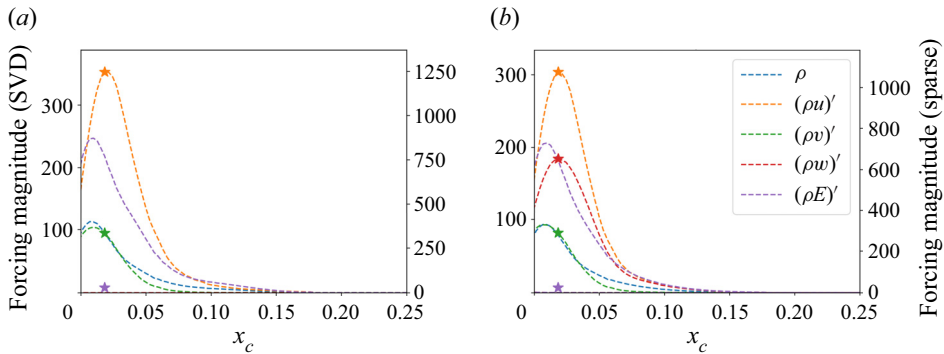


Figure 14. The surface forcing on the suction side of the aerofoil at $St = 5.22$ for $\beta = 0$, and $St = 5.90$ for $\beta = 20\pi$. The sparse components are shown with a star. (a) Surface forcing on the suction side of the aerofoil for $\beta = 0$. (b) Surface forcing on the suction side of the aerofoil for $\beta = 20\pi$.

located near the maximum value of the full-SVD forcing mode. Whilst this may show that choosing the largest value of the forcing mode is a good candidate for the sparse forcing mode, we emphasise that the optimisation procedure is not biased by any knowledge of the full forcing mode, and all physical mechanisms and spatial locations are weighted equally.

Whilst forcing in the shear region may provide the optimal response, it is rather impractical for flow actuation purposes. Therefore, we conclude this section by using the windowing matrix \mathbf{B} to conduct our analysis on the surface of the aerofoil. Figure 14 shows the forcing distribution along the suction side of the aerofoil as a function of the distance along the chord x_c . In the full resolvent analysis, most of the forcing is concentrated near the leading edge of the aerofoil, agreeing with the non-windowed case. In both spanwise wavenumber cases, the sparse mode is once again a single-point momentum-based forcing and is located at the maximum value for the kinetic energy of the full forcing mode. This provides the optimal compromise between forcing with $(\rho u)'$ at its maximum value and $(\rho v)'$ at its maximum, which in the full case is located to the left of $(\rho u)'$. Even though there is a $(\rho E)'$ -component, this is simply a consequence of the kinetic part of the energy, since $(\rho E)' = \rho' \|\mathbf{u}_0\|^2/2 + \rho_0 \mathbf{u}' \cdot \mathbf{u}_0 + P' / (\gamma - 1)$, and there is no thermodynamic contribution to the linearised total energy. The importance of grouping momentum together into one coherent strategy is evident from figure 14, as the directional information of the actuation is crucial in both the full and sparse resolvent analyses to achieve the optimal gain. This information would otherwise be lost. Moreover, by grouping the momentum together, we strike a compromise over choosing a forcing that is optimal for each isolated velocity component.

5. Conclusion

By reformulating an optimal-input analysis as a Riemannian optimisation problem, we are able to tailor a resolvent analysis to uncover sparse forcing modes and their corresponding response modes. By designing a cost functional based on the ratio of the energy gain to the L_1 -norm of the forcing, we are able to find forcing modes that provide the largest gain whilst being spatially sparse. To test the method within the context of resolvent analyses performed on steady base-flows and time-averaged mean-flows, we considered two flow examples: plane Poiseuille flow in the linearly stable regime, and the turbulent flow past an aerofoil.

For plane Poiseuille flow, two forcing frequencies were considered. At the first frequency, located at the maximum gain of the full-SVD analysis, the sparse forcing mode consisted of an $\alpha = 1$ stripe at a single y -location, just above the critical layer. Conversely, for the second frequency located at the second peak of the full analysis, the forcing consisted of a single spatial forcing near the $\alpha = 2$ critical layer. For the first case, utilising an $\alpha = 1$ forcing is critical in obtaining a high gain, therefore the optimisation sparsifies the forcing only in the y -direction. However, in the second case, there is a much lower separation in the effectiveness of different forcing mechanisms, as indicated by the ratio of the singular values, meaning that the sparse procedure is able to sparsify further whilst still providing a large gain.

In the turbulent flow past an aerofoil, all sparse modes consisted of single spatial locations, with the sparsification procedure also identifying momentum-based forcing as the optimum physical mechanism. For two different spanwise wavenumbers, an analysis of the resolvent wavemakers shows that forcing in the shear layer provides the optimal gain, with the sparse procedure focusing on the location of the maximal value of the full forcing modes. To identify an implementable actuator position, we also considered a windowed analysis where the forcing modes are confined to the surface of the aerofoil. Again, we achieve single-point momentum-based actuation positions that are found to be a compromise among the optimal locations for each independent velocity component. This emphasises the importance of designing an appropriate vector for the L_1 -norm, as the directional information would have been lost had we not grouped momentum together.

Overall, the sparse optimisation procedure provides an unbiased optimal sparsification of the flow and is able to adapt to the different forcing strategies available at different frequencies. Although the aerofoil results show that choosing the maxima of the SVD is a good candidate for a sparse forcing vector, the plane Poiseuille example shows that both single-point and multi-point forcing modes can be found, depending on the low-rank nature and physical mechanisms furnished by the resolvent. Based on our results, it can be postulated that in more complex systems, such as those stemming from aeroacoustic or combustion problems, where multiple physical mechanisms are at play, the sparse resolvent would be able to adapt to the optimal physical mechanisms present at each frequency, and would even be able to combine in an optimal sparse way these different mechanisms in order to achieve the largest gain. Investigating the sparse optimisation procedure on these types of flows would therefore be an interesting future direction of study. Furthermore, as the choice of cost functional for the purpose of sparsification is not unique, the design of other functionals, such as those that could allow for a tuning of sparsity versus gain, provides another area for future investigation.

Although we have not considered them in our study, we note further that recent efforts have been made to extend resolvent analysis both to periodic flows (Padovan, Otto &

Rowley 2020) and also to the nonlinear regime (Rigas, Sipp & Colonius 2021). As the techniques that we have presented carry over to both these cases without significant modification, these present interesting avenues for future investigations. Moreover, the usefulness of Riemannian optimisation in tailoring input–output analyses to specific flow applications is not limited to our sparse analysis. As well as being able to design cost functionals in order to uncover different aspects of the resolvent, the manifold to which we confine the forcing modes can be changed. The result is a rich landscape of possibilities in which resolvent analyses can be extended, with the traditional SVD-based approach being just one such choice.

Acknowledgements. We would like to thank J.H.M. Ribeiro, T.R. Ricciardi and D. House for fruitful discussions on resolvent analysis. We are also grateful for the insightful comments of the referees in the peer review process.

Funding. K.T. and P.J.S. acknowledge the support from the Army Research Office (grant W911NF-21-1-0060, program officer M.J. Munson). K.T. also acknowledges the support from the Office of Naval Research (grant N0014-19-1-2460, program officer D.R. Gonzalez).

Declaration of interests. The authors report no conflict of interest.

Author ORCIDs.

-  Calum S. Skene <https://orcid.org/0000-0003-0994-2013>;
-  Chi-An Yeh <https://orcid.org/0000-0003-0426-8381>;
-  Peter J. Schmid <https://orcid.org/0000-0002-2257-8490>;
-  Kunihiko Taira <https://orcid.org/0000-0002-3762-8075>.

Appendix. Convergence results

In this appendix, we consider the numerical details of the optimisation procedure. In order to provide an overview, we will present the results stemming from the plane Poiseuille example at $\omega = 0.278$ considered in § 4, which is representative of all cases considered in this paper.

Figure 15 shows how the results of the optimisation procedure depend on δ . We can see from figures 15(a) and 15(b) that there is an initial region in which the optimisation procedure has a large L_1 -norm and that the gain is in line with that obtained from the SVD. After $\delta = 2^{-5}$, the pseudo-Huber-norm starts to behave more like the L_1 -norm, as shown in figures 15(b) and 15(d), and both the gain and the L_1 -norm of the forcing decrease. At $\delta = 5^{-7}$, the pseudo-Huber-norm and the L_1 -norm have converged, and we can stop the optimisation procedure. These figures also highlight the strong dependence of the gain on the L_1 -norm of the forcing for these examples, with the gain decreasing almost exactly in line with the L_1 -norm.

Figures 16(a) and 16(b) show the norm of the gradient provided by the optimisation procedure as a function of the number of iterations. For the case $\omega = 0.278$, the number of iterations is in line with the work of Foures *et al.* (2013), with perhaps a few more iterations needed in our case. The number of iterations required, as well as the non-smoothness of the gradient norms, is indicative of the difficulty of the gradient-based optimisation. This is especially highlighted by figure 16(b), where the optimisation terminates earlier due to a stagnation of the cost functional. This showcases the importance of using both a relaxation parameter δ as well as an optimisation procedure such as the conjugate gradient algorithm in order to achieve converged results.

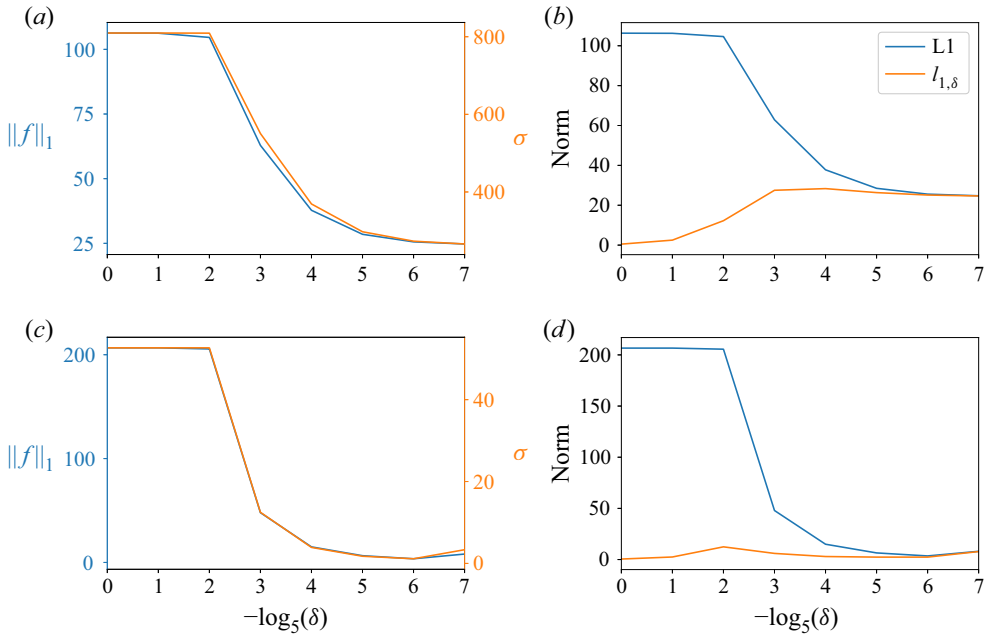


Figure 15. Dependence of the results on δ for the plane Poiseuille flow examples. (a) L_1 -norm and gains with δ for $\omega = 0.278$. (b) Convergence of the pseudo-Huber-norm to the true L_1 -norm for $\omega = 0.278$. (c) L_1 -norm and gains with δ for $\omega = 1.14$. (d) Convergence of the pseudo-Huber-norm to the true L_1 -norm for $\omega = 1.14$.




Figure 16. The optimisation convergence behaviour for the plane Poiseuille flow examples. Convergence behaviour of the conjugate gradient algorithm for plane Poiseuille flow at (a) $\omega = 0.278$ and $\delta = 5^{-7}$, (b) $\omega = 1.14$ and $\delta = 5^{-7}$.

REFERENCES

ABSIL, P.-A., MAHONY, R. & SEPULCHRE, R. 2007 *Optimization Algorithms on Matrix Manifolds*. Princeton University Press.

AMESTOY, P.R., BUTTARI, A., L'EXCELLENT, J.-Y. & MARY, T. 2019 Performance and scalability of the block low-rank multifrontal factorization on multicore architectures. *ACM Trans. Math. Softw.* **45**, 1–26.

AMESTOY, P.R., DUFF, I.S., KOSTER, J. & L'EXCELLENT, J.-Y. 2001 A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM J. Matrix Anal. Applics.* **23** (1), 15–41.

BALAY, S. *et al.* 2021a PETSc users manual. *Tech. Rep.* ANL-95/11 – Revision 3.15. Argonne National Laboratory.

BALAY, S. *et al.* 2021b PETSc Web page. <https://petsc.org/>.

- BALAY, S., GROPP, W.D., MCINNES, L.C. & SMITH, B.F. 1997 Efficient management of parallelism in object oriented numerical software libraries. In *Modern Software Tools in Scientific Computing* (ed. E. Arge, A.M. Bruaset & H.P. Langtangen), pp. 163–202. Birkhauser.
- BOUMAL, N. & ABSIL, P.-A. 2015 Low-rank matrix completion via preconditioned optimization on the Grassmann manifold. *Linear Algebr. Applics.* **475**, 200–239.
- BOUMAL, N., MISHRA, B., ABSIL, P.-A. & SEPULCHRE, R. 2014 Manopt, a Matlab toolbox for optimization on manifolds. *J. Mach. Learn Res.* **15**, 1455–1459.
- BRÈS, G.A., HAM, F.E., NICHOLS, J.W. & LELE, S.K. 2017 Unstructured large-eddy simulations of supersonic jets. *AIAA J.* **55** (4), 1164–1184.
- BUBE, K.P. & LANGAN, R.T. 1997 Hybrid ℓ_1/ℓ_2 minimization with applications to tomography. *Geophysics* **62** (4), 1183–1195.
- BUBE, K.P. & NEMETH, T. 2007 Fast line searches for the robust solution of linear systems in the hybrid ℓ_1/ℓ_2 and Huber norms. *Geophysics* **72** (2), A13–A17.
- CATTAFESTA, L.N. & SHEPLAK, M. 2011 Actuators for active flow control. *Annu. Rev. Fluid Mech.* **43**, 247–272.
- CHU, B.T. 1965 On the energy transfer to small disturbances in fluid flow (part I). *Acta Mechanica* **1** (3), 215–234.
- DALCIN, L.D., PAZ, R.R., KLER, P.A. & COSIMO, A. 2011 Parallel distributed computing using Python. *Adv. Water Resour.* **34** (9), 1124–1139, New Computational Methods and Software Tools.
- DOUGLAS, S.C., AMARI, S. & KUNG, S.-Y. 1998 Gradient adaptation under unit-norm constraints. In *Ninth IEEE Signal Processing Workshop on Statistical Signal and Array Processing (Cat. No.98TH8381)*, pp. 144–147. IEEE.
- FARRELL, B.F. & IOANNOU, P.J. 1993 Stochastic forcing of the linearized Navier–Stokes equations. *Phys. Fluids A: Fluid Dyn.* **5** (11), 2600–2609.
- FOSAS DE PANDO, M. 2020 IBMOS: immersed boundary method optimization and stability. <https://doi.org/10.5281/zenodo.3757783>.
- FOSAS DE PANDO, M. & SCHMID, P.J. 2017 Optimal frequency-response sensitivity of compressible flow over roughness elements. *J. Turbul.* **18** (4), 338–351.
- FOSAS DE PANDO, M., SCHMID, P.J. & LELE, S.K. 2014 Parametric sensitivity for large-scale aeroacoustic flows. In *Proceedings of the 2014 Summer Program*, pp. 365–374. Center for Turbulence Research, Stanford University.
- FOURES, D.P.G., CAULFIELD, C.P. & SCHMID, P.J. 2013 Localization of flow structures using ∞ -norm optimization. *J. Fluid Mech.* **729**, 672–701.
- GIANNETTI, F. & LUCHINI, P. 2007 Structural sensitivity of the first instability of the cylinder wake. *J. Fluid Mech.* **581**, 167–197.
- HERNANDEZ, V., ROMAN, J.E. & VIDAL, V. 2005 SLEPc: a scalable and flexible toolkit for the solution of eigenvalue problems. *ACM Trans. Math. Softw.* **31** (3), 351–362.
- HUANG, W., ABSIL, P.-A. & GALLIVAN, K.A. 2016 A Riemannian BFGS method for nonconvex optimization problems. In *Numerical Mathematics and Advanced Applications ENUMATH 2015* (ed. B. Karasözen, M. Manguoğlu, M. Tezer-Sezgin, S. Göktepe & Ö. Uğur), Lecture Notes in Computational Science and Engineering, vol. 112, pp. 627–634. Springer.
- JEUN, J., NICHOLS, J.W. & JOVANOVIĆ, M.R. 2016 Input–output analysis of high-speed axisymmetric isothermal jet noise. *Phys. Fluids* **28** (4), 047101.
- JOVANOVIĆ, M.R. 2004 Modeling, analysis, and control of spatially distributed systems. PhD thesis, University of California, Santa Barbara.
- JOVANOVIĆ, M.R. & BAMIEH, B. 2005 Componentwise energy amplification in channel flows. *J. Fluid Mech.* **534**, 145–183.
- KHALIGHI, Y., NICHOLS, J.W., HAM, F., LELE, S.K. & MOIN, P. 2011 Unstructured large eddy simulation for prediction of noise issued from turbulent jets in various configurations. *17th AIAA/CEAS Aeroacoustics Conference (32nd AIAA Aeroacoustics Conference)*. *AIAA Paper* 2011-2886.
- LIU, Q., SUN, Y., YEH, C.-A., UKEILEY, L.S., CATTAFESTA, L.N. & TAIRA, K. 2021 Unsteady control of supersonic turbulent cavity flow based on resolvent analysis. *J. Fluid Mech.* **925**, A5.
- LUHAR, M., SHARMA, A.S. & MCKEON, B.J. 2014 Opposition control within the resolvent analysis framework. *J. Fluid Mech.* **749**, 597–626.
- MCKEON, B.J. & SHARMA, A.S. 2010 A critical-layer framework for turbulent pipe flow. *J. Fluid Mech.* **658**, 336–382.
- PADOVAN, A., OTTO, S.E. & ROWLEY, C.W. 2020 Analysis of amplification mechanisms and cross-frequency interactions in nonlinear flows via the harmonic resolvent. *J. Fluid Mech.* **900**, A14.

- PALADINI, E., BENEDDINE, S., DANDOIS, J., SIPP, D. & ROBINET, J.-C. 2019 Transonic buffet instability: from two-dimensional airfoils to three-dimensional swept wings. *Phys. Rev. Fluids* **4**, 103906.
- PRINGLE, C.C.T., WILLIS, A.P. & KERSWELL, R.R. 2012 Minimal seeds for shear flow turbulence: using nonlinear transient growth to touch the edge of chaos. *J. Fluid Mech.* **702**, 415–443.
- QADRI, U.A. & SCHMID, P.J. 2017 Frequency selection mechanisms in the flow of a laminar boundary layer over a shallow cavity. *Phys. Rev. Fluids* **2**, 013902.
- RIBEIRO, J.H.M., YEH, C.-A. & TAIRA, K. 2020 Randomized resolvent analysis. *Phys. Rev. Fluids* **5**, 033902.
- RIGAS, G., SIPP, D. & COLONIUS, T. 2021 Nonlinear input/output analysis: application to boundary layer transition. *J. Fluid Mech.* **911**, A15.
- ROSENBERG, K., SYMON, S. & MCKEON, B.J. 2019 Role of parasitic modes in nonlinear closure via the resolvent feedback loop. *Phys. Rev. Fluids* **4**, 052601.
- SCHMID, P.J. & BRANDT, L. 2014 Analysis of fluid systems: stability, receptivity, sensitivity: lecture notes from the FLOW-NORDITA summer school on advanced instability methods for complex flows, Stockholm, Sweden, 2013. *Appl. Mech. Rev.* **66** (2), 024803.
- SCHMID, P.J. & HENNINGSON, D.S. 2001 *Stability and Transition in Shear Flows*, Applied Mathematical Sciences, vol. 142. Springer.
- SKENE, C.S., RIBEIRO, J.H.M. & TAIRA, K. 2022 csskene/linear-analysis-tools: initial release. <https://doi.org/10.5281/zenodo.6550726>.
- SUN, Y., TAIRA, K., CATTAFESTA, L.N. & UKEILEY, L.S. 2017 Biglobal instabilities of compressible open-cavity flows. *J. Fluid Mech.* **826**, 270–301.
- TAIRA, K., BRUNTON, S.L., DAWSON, S.T.M., ROWLEY, C.W., COLONIUS, T., MCKEON, B.J., SCHMIDT, O.T., GORDEYEV, S., THEOFILIS, V. & UKEILEY, L.S. 2017 Modal analysis of fluid flows: an overview. *AIAA J.* **55** (12), 4013–4041.
- TAIRA, K. & COLONIUS, T. 2007 The immersed boundary method: a projection approach. *J. Comput. Phys.* **225** (2), 2118–2137.
- TAIRA, K., HEMATI, M.S., BRUNTON, S.L., SUN, Y., DURAISAMY, K., BAGHERI, S., DAWSON, S. & YEH, C.-A. 2020 Modal analysis of fluid flows: applications and outlook. *AIAA J.* **58** (3), 998–1022.
- TOEDTLI, S.S., LUHAR, M. & MCKEON, B.J. 2019 Predicting the response of turbulent channel flow to varying-phase opposition control: resolvent analysis as a tool for flow control design. *Phys. Rev. Fluids* **4**, 073905.
- TOWNSEND, J., KOEP, N. & WEICHWALD, S. 2016 Pymanopt: a Python toolbox for optimization on manifolds using automatic differentiation. *J. Mach. Learn. Res.* **17** (137), 1–5.
- TREFETHEN, L.N., TREFETHEN, A.E., REDDY, S.C. & DRISCOLL, T.A. 1993 Hydrodynamic stability without eigenvalues. *Science* **261** (5121), 578–584.
- VREMAN, A.W. 2004 An eddy-viscosity subgrid-scale model for turbulent shear flow: algebraic theory and applications. *Phys. Fluids* **16** (10), 3670–3681.
- YEH, C.-A., BENTON, S.I., TAIRA, K. & GARMANN, D.J. 2020 Resolvent analysis of an airfoil laminar separation bubble at $Re = 500\,000$. *Phys. Rev. Fluids* **5** (8), 083906.
- YEH, C.-A. & TAIRA, K. 2019 Resolvent-analysis-based design of airfoil separation control. *J. Fluid Mech.* **867**, 572–610.