



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/188910/>

Version: Accepted Version

Article:

Skene, CS, Yeh, C-A, Schmid, PJ et al. (2022) Sparsifying the resolvent forcing mode via gradient-based optimisation. *Journal of Fluid Mechanics*, 944. A52. ISSN: 0022-1120

<https://doi.org/10.1017/jfm.2022.519>

This article has been published in a revised form in *Journal of Fluid Mechanics* [<http://doi.org/10.1017/jfm.2022.519>]. This version is free to view and download for private research and study only. Not for re-distribution, re-sale or use in derivative works. © The Author(s), 2022. Published by Cambridge University Press.

Reuse

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs (CC BY-NC-ND) licence. This licence only allows you to download this work and share it with others as long as you credit the authors, but you can't change the article in any way or use it commercially. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Banner appropriate to article type will appear here in typeset article

Sparsifying the resolvent forcing mode via gradient-based optimisation

Calum S. Skene¹†‡, Chi-An Yeh¹¶, Peter J. Schmid², and Kunihiro Taira¹

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

²Department of Mathematics, Imperial College London, London SW7 2AZ, United Kingdom

(Received xx; revised xx; accepted xx)

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 a Reynolds number of $Re = 4000$ and turbulent flow past an NACA0012 aerofoil at $Re = 23000$. 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.

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 (Trefethen *et al.* 1993; Farrell & Ioannou 1993). By sweeping through frequencies, structural mechanisms that provide efficient means of flow amplification, as well as 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).

† 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

34 With the resolvent analysis being able to reveal the input-output relationship
 35 with respect to the chosen base state (Jovanović & Bamieh 2005), it naturally
 36 serves as a valuable tool to design flow control techniques. Past studies including
 37 Luhar *et al.* (2014), Yeh & Taira (2019), Toedtli *et al.* (2019), and Liu *et al.* (2021)
 38 have demonstrated that physical insights revealed from resolvent analysis provide
 39 valuable guidance of developing effective and efficient actuation strategies.

40 Traditionally, modal analysis techniques for fluid flows (Taira *et al.* 2017,
 41 2020) have been founded on L_2 -based norms, which can lead to global spatial
 42 structures. For the resolvent analysis, this translates to having forcing modes
 43 that are spatially supported over a large region. It should however be realised
 44 that actuation inputs cannot be applied over a large spatial region in practical
 45 engineering flow control settings. In general, flow control inputs can only be intro-
 46 duced as localised actuation inputs. To address this point, we consider sparsity-
 47 promoting approaches to specifically target resolvent forcing modes that have
 48 spatially compact support, i.e., are spatially sparse. We also note that sparsity-
 49 promoting techniques may also help identify appropriate variables through which
 50 control inputs can be added to the flow for enhanced amplification. This piece
 51 of information is important in selecting the appropriate type of actuators to
 52 introduce control input to the flow field (Cattafesta & Sheplak 2011).

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

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

71 Following this approach, our study considers resolvent analysis as an opti-
 72 misation problem where forcing modes are sought that maximise a prescribed
 73 cost functional. In order to obtain spatially sparse forcing modes, we propose
 74 a gradient-based algorithm that maximises the energy of the output whilst
 75 minimising the L_1 -norm of the forcing, which is also constrained to have unit
 76 energy. To provide an initial assessment of our proposed method we consider two
 77 examples. Firstly, we consider the same plane Poiseuille setup as in Foures *et al.*
 78 (2013), allowing us to qualitatively assess the differences between localisation
 79 strategies for initial conditions and for forced problems. Secondly, we consider
 80 turbulent flow past an aerofoil using the same mean-flow as Yeh & Taira (2019),
 81 providing an assessment of the method in a higher Reynolds number, turbulent
 82 flow.

83 The structure of the paper is as follows. Section 2 outlines the mathematical

84 formulation of the paper and contains an introduction to the resolvent operator,
 85 a background on Riemannian optimisation and how we utilise it to find optimal,
 86 sparse resolvent modes, and a discussion of wavemakers in the context of a
 87 resolvent analysis. In section 3 we discuss the numerical setup, with the results
 88 subsequently being presented in section 4. Conclusions are offered in section 5.

89 2. Formulation

90 2.1. The resolvent operator

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

$$93 \quad \mathbf{G} \frac{d\mathbf{q}}{dt} = \mathcal{N}(\mathbf{q}), \quad (2.1)$$

94 where \mathbf{q} is the spatially discretized state vector, and \mathcal{N} represents the right-
 95 hand side of the unforced Navier–Stokes equations. Including the mass-matrix \mathbf{G}
 96 in equation (2.1) means that this form could represent either the compressible
 97 Navier–Stokes equations or the incompressible equations where there is no time-
 98 derivative in the continuity equation. By linearising this equation about a base
 99 flow \mathbf{q}_0 , we can write the system in input-output form as

$$100 \quad \mathbf{G} \frac{d\mathbf{q}}{dt} = \mathbf{L}_{\mathbf{q}_0} \mathbf{q} + \mathbf{B} \mathbf{f}, \quad (2.2)$$

$$101 \quad \mathbf{y} = \mathbf{C} \mathbf{q}, \quad (2.3)$$

102 where $\mathbf{L}_{\mathbf{q}_0}$ is the linearised Navier–Stokes operator (Jeun *et al.* 2016). The matrix
 103 \mathbf{B} allows for the introduced forcing \mathbf{f} (input) to be windowed in space or restricted
 104 to specific equations or state variables. In an analogous manner, the matrix \mathbf{C}
 105 allows for a similar windowing to be applied to the output \mathbf{y} .

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

$$109 \quad \hat{\mathbf{y}} = \mathbf{C}(-i\omega \mathbf{G} - \mathbf{L}_{\mathbf{q}_0})^{-1} \mathbf{B} \hat{\mathbf{f}}. \quad (2.4)$$

110 Through this equation, the resolvent operator is defined via $\mathcal{H}_{\mathbf{q}_0} \equiv \mathbf{C}(-i\omega \mathbf{G} -$
 111 $\mathbf{L}_{\mathbf{q}_0})^{-1} \mathbf{B}$. The form of equation (2.4) shows that the resolvent operator is equiva-
 112 lent to a transfer function which maps the forcing to its time-asymptotic response.
 113 Before we discuss the meaning of the resolvent in fluid dynamics it is worth
 114 considering the Laplace variable ω . If the operator $\mathbf{L}_{\mathbf{q}_0}$ is stable then ω is real and
 115 (2.4) is obtainable via the Fourier transform. However, if $\mathbf{L}_{\mathbf{q}_0}$ is unstable then more
 116 care is needed. Indeed, for unstable $\mathbf{L}_{\mathbf{q}_0}$ the time-asymptotic response is not given
 117 via (2.4) and is instead a combination of the exponentially growing disturbance
 118 given by the most unstable eigenvector and the forced response given by the
 119 resolvent. In order to separate these two mechanisms a complex value for ω can
 120 be used leading to the concept of a time-discounted resolvent analysis (Jovanović
 121 2004). Choosing complex values for ω means that the imaginary part can be
 122 chosen such that the forced response ‘rises above’ the exponentially growing
 123 disturbance due to the unstable nature of $\mathbf{L}_{\mathbf{q}_0}$, allowing for the forced dynamics
 124 to be probed (see Yeh *et al.* (2020) for more details).

125 In the context of fluid dynamics, the resolvent can be interpreted in two main
 126 ways. Firstly, choosing \mathbf{q}_0 to be a steady solution to the unforced Navier–Stokes
 127 equations leads to a non-normal stability study of the flow. In this manner, the

128 resolvent identifies forcing structures that lead to particularly efficient amplifi-
 129 cation in the dynamics despite the stable nature of the flow (Trefethen *et al.*
 130 1993). Secondly, using a time-averaged mean-flow for \mathbf{q}_0 leads to the resolvent
 131 formulation of turbulence (McKeon & Sharma 2010). The resolvent in this
 132 instance can be used to identify the coherent structures that arise via disturbances
 133 caused by the non-linear terms.

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

$$141 \quad \mathbf{f}_{\text{opt}} = \arg \max_{\mathbf{f}} \frac{\|\mathbf{y}\|_{\mathbf{w}_q}}{\|\mathbf{f}\|_{\mathbf{w}_f}}, \quad (2.5)$$

142 where the norms are defined as $\|\mathbf{a}\|_{\mathbf{W}}^2 = \mathbf{a}^H \mathbf{W} \mathbf{a}$ with \mathbf{W} being a positive definite
 143 weight matrix. We allow for the weight matrix for the forcing (\mathbf{W}_f) and response
 144 (\mathbf{W}_q) to be different. These matrices are problem dependent, and are chosen so
 145 that the norms represent appropriate measures of the energy (see sections 3.1 and
 146 3.2 for examples). The cost functional in this case is the gain. To link the weighted
 147 norms to the two norm, it is useful to also consider the Cholesky decomposition
 148 $\mathbf{W} = \mathbf{M}^H \mathbf{M}$. The optimal forcing has the corresponding output $\tilde{\mathbf{y}}_{\text{opt}} = \mathcal{H} \mathbf{f}_{\text{opt}}$ with
 149 the amount of amplification being measured by the gain $\sigma = \|\tilde{\mathbf{y}}_{\text{opt}}\|_{\mathbf{w}_q} / \|\mathbf{f}_{\text{opt}}\|_{\mathbf{w}_f}$.
 150 This problem can be solved by taking the SVD of $\mathbf{M}_q \mathcal{H} \mathbf{M}_f^{-1}$, whose maximum
 151 singular triplet is $(\sigma, \mathbf{M}_q \mathbf{y}_{\text{opt}}, \mathbf{M}_f \mathbf{f}_{\text{opt}})$.

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

160 2.2. Sparsification via Riemannian optimisation

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

$$164 \quad \sigma = \frac{\|\mathcal{H} \mathbf{f}\|_{\mathbf{w}_q}^2}{\|\mathbf{f}\|_{\mathbf{w}_f}^2}. \quad (2.6)$$

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

$$168 \quad \sigma = \|\mathcal{H} \mathbf{f}\|_{\mathbf{w}_q}^2, \quad (2.7)$$

169 where the forcing is confined to the manifold given by the constraint $\|\mathbf{f}\|_{\mathbf{w}_f}^2 = 1$.
 170 This is an equivalent problem to (2.6) because the resolvent is linear and hence




Figure 1: A sketch illustrating the concept of Riemannian optimisation. First, the Euclidean gradient $\nabla_{\mathbf{f}_M^n} \sigma^{(L_2)}$ is found from the vector \mathbf{f}_M^n that is situated on the hypersphere \mathcal{S} . This vector is then mapped to the tangent space $T_{\mathbf{f}_M^n} \mathcal{S}$ via the projection $\text{Proj}_{\mathbf{f}_M^n}$. 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 $R_{\mathbf{f}_M^n}$, yielding the updated forcing vector \mathbf{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.

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

175 Whilst we could conduct an unconstrained optimisation by enforcing $\|\mathbf{f}\|_{\mathbf{w}_f}^2 = 1$
 176 with a Lagrange multiplier (Pringle *et al.* 2012), we instead take account of this
 177 constraint directly in the update step. This results in a similar approach to that
 178 of Foures *et al.* (2013), where a geometric approach was used to ensure that the
 179 unit-norm condition is satisfied when stepping in the search direction. In general,
 180 carrying out an optimisation where the input is constrained to a manifold is
 181 known as Riemannian optimisation (Absil *et al.* 2007).

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

$$184 \quad \sigma^{(L_2)} = \|\mathbf{M}_q \mathcal{H} \mathbf{M}_f^{-1} \mathbf{M}_f \mathbf{f}\|_2, \quad (2.8)$$

185 subject to $\|\mathbf{M}_f \mathbf{f}\|_2 = 1$. By expressing the problem in this form, we have
 186 reformulated the problem in terms of the L_2 -norm, and hence we are optimising
 187 with respect to the vector $\mathbf{f}_M = \mathbf{M}_f \mathbf{f}$, which we constrain to have unit L_2 -
 188 norm. The manifold for this problem then becomes the complex-hypersphere
 189 $\mathcal{S} = \{\mathbf{y} \mid \mathbf{y}^H \mathbf{y} = 1\}$.

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

191 gradient

$$192 \quad \nabla_{\mathbf{f}_M} \sigma^{(L_2)} = \frac{\partial \sigma^{(L_2)}}{\partial \mathbf{f}_M}. \quad (2.9)$$

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

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

$$206 \quad \text{grad } \sigma^{(L_2)}(\mathbf{f}_M) = (1 - \mathbf{f}_M^H \mathbf{f}_M) \nabla_{\mathbf{f}_M} \sigma^{(L_2)} = \text{Proj}_{\mathbf{f}_M} (\nabla_{\mathbf{f}_M} \sigma^{(L_2)}), \quad (2.10)$$

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

209 Now that we have defined appropriate gradients, we must also define how to
 210 step in the direction of steepest ascent. For the unconstrained optimisation, we
 211 may simply add a scalar multiple (the step size) of this gradient onto our current
 212 value of the forcing. However, for the Riemannian optimisation, this will result
 213 in a vector that no longer falls on the manifold, as noted above. The equivalent
 214 procedure in this case is the notion of a retraction. A retraction is a mapping
 215 $R_{\mathbf{x}}(\boldsymbol{\xi}) : \mathbb{T}_{\mathbf{x}}\mathcal{S} \rightarrow \mathcal{S}$ such that $R_{\mathbf{x}}(\mathbf{0}) = \mathbf{x}$ and $DR_{\mathbf{x}}(\mathbf{0}) = \text{id}_{\mathbb{T}_{\mathbf{x}}\mathcal{S}}$. In other words,
 216 a retraction maps vectors tangent to the manifold at \mathbf{x} to other vectors on the
 217 manifold such that for $\boldsymbol{\xi} = \mathbf{0}$ it maps \mathbf{x} to itself, and that the derivative of the
 218 mapping at $\boldsymbol{\xi} = \mathbf{0}$ is the identity. For the hypersphere, we have the retraction

$$219 \quad R_{\mathbf{x}}(\boldsymbol{\xi}) = \frac{\mathbf{x} + \boldsymbol{\xi}}{\|\mathbf{x} + \boldsymbol{\xi}\|}. \quad (2.11)$$

220 Once the gradient is found, we can then update the forcing using the map
 221 $R_{\mathbf{f}_M}(\alpha \text{grad } \sigma^{(L_2)}(\mathbf{f}_M))$, where α denotes the step size. By writing $\cos(\theta) =$
 222 $1/\sqrt{1 + \alpha^2}$, we can also express the update step as

$$223 \quad \mathbf{f}_M^{n+1} = R_{\mathbf{f}_M}(\alpha \text{grad } \sigma^{(L_2)}(\mathbf{f}_M^n)) = \cos(\theta) \mathbf{f}_M^n + \sin(\theta) \text{grad } \sigma^{(L_2)}(\mathbf{f}_M^n), \quad (2.12)$$

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

230 The main advantage of phrasing the optimal forcing-output problem in this
 231 way is its generality. Whilst we have shown how we can obtain the same result as
 232 the SVD (and it is actually possible to get the higher-order singular values in this
 233 manner by considering a different manifold), we are free to change how we define

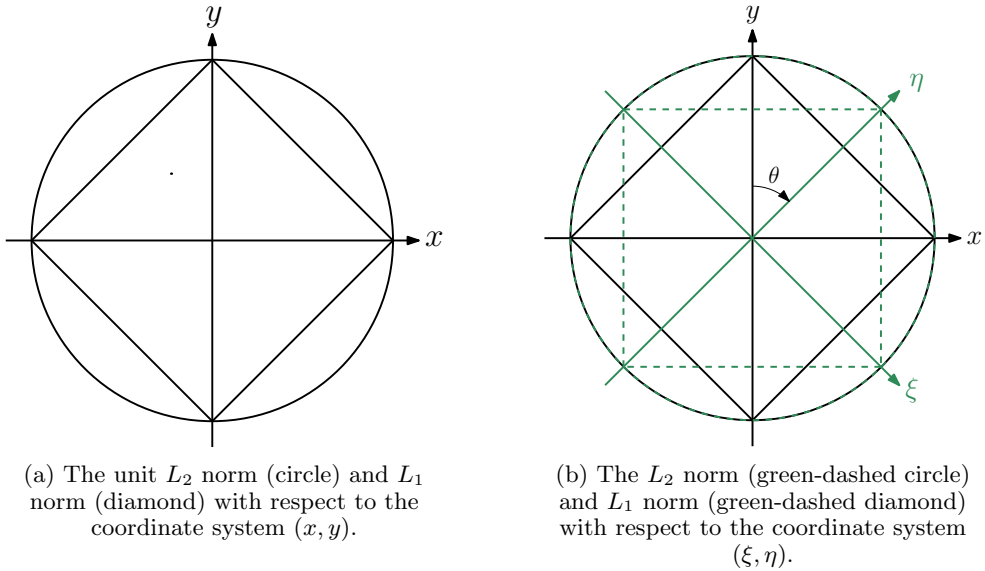


Figure 2: Sketches of the L_2 - and L_1 -norms. The effect of a coordinate rotation on the norms is also demonstrated.

234 the gain. The SVD can only find the gain in the L_2 -norm sense. This means that
 235 the input is measured by an energy norm, leading to the global structures seen
 236 in many studies. In order to introduce sparsification, we consider the use of the
 237 1-norm.

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

248 One important consideration when using an alternative norm, such as the L_1 -
 249 norm, is illustrated by figure 2b. Here we have again shown the unit L_2 - and
 250 L_1 -norms, but this time for the coordinate system (η, ξ) which is obtained via
 251 a rotation of the coordinate system (x, y) by an angle θ . In this new coordi-
 252 nate system, the L_2 -norm still represents a circle, which is invariant under this
 253 transformation. However, the unit L_1 -norm is affected, and its square shape is
 254 rotated by the angle θ . This means that in this new coordinate system the sparse
 255 vectors, where the square touches the circle, are different to those of the original
 256 coordinate system (x, y) . In other words, what is considered sparse is completely
 257 defined by how we choose to represent our vectors. In practise, we must be careful
 258 when choosing the vector of which we take the L_1 -norm. We therefore choose
 259 to take the L_1 -norm of a vector that leaves the L_2 -norm unchanged, yet has
 260 appropriate axes for best defining the sparsity we aim to achieve. In terms of

261 resolvent analysis, this transformation is used to maintain the physical relevance
 262 of the sparsification. Specific examples are described in sections 3.1 and 3.2.

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

$$265 \quad \sigma^{(L_1)} = \frac{\sigma^{(L_2)}}{\|\mathbf{T}(\mathbf{f}_M)\|_1} = \frac{\|\mathbf{M}_q \mathcal{H} \mathbf{M}_f^{-1} \mathbf{f}_M\|_2}{\|\mathbf{T}(\mathbf{f}_M)\|_1}, \quad (2.13)$$

266 still subject to the forcing \mathbf{f}_M having a unit-energy norm. The transformation
 267 \mathbf{T} in the denominator is a transformation of the vector \mathbf{f}_M to another vector.
 268 Hence, the vector in the denominator need not be equal to the forcing vector
 269 \mathbf{f}_M as, based on the discussion of the previous paragraph, this may not be
 270 physically relevant. However, by ensuring $\|\mathbf{T}(\mathbf{f}_M)\|_2 = \|\mathbf{f}_M\|_2$, we maintain the
 271 geometric interpretation of sparsity illustrated by figure 2a, albeit in a much
 272 higher dimensional space. By dividing the usual gain by the 1-norm of the vector
 273 $\mathbf{T}(\mathbf{f}_M)$, we are in effect promoting sparsity, with sparsity defined as a vector
 274 $\mathbf{T}(\mathbf{f}_M)$ with a minimal number of non-zero entries. Optimising the gain in this
 275 form will seek a compromise between providing a large gain in energy whilst
 276 ensuring the spatial sparsity of the forcing. Indeed, the maximal nature of $\sigma^{(L_1)}$
 277 means that obtaining a response with larger energy requires a forcing structure
 278 that is less sparse. Likewise, making the forcing more sparse leads to a less
 279 energetic response.

280 As in the study of Foures *et al.* (2013), who considered a similar optimisation
 281 problem for localising flow structures obtained in transient growth studies, our
 282 cost functional is non-convex. This means that any solution to the optimisation
 283 problem is only guaranteed to be a local, rather than global, maximum of the
 284 cost functional. In the case of transient growth this led to multiple branches
 285 of solutions being found during the optimisation, which could be discovered
 286 by running the problem with multiple starting guesses for the gradient-based
 287 optimisation. However, despite running multiple instances of each optimisation
 288 with different initial guesses in our following examples, no differences in the
 289 solution could be found apart from symmetries of the flow which are to be
 290 expected. Whilst this does not confirm that our results are truly the global
 291 optimum, it does highlight a difference between localising forcings for driven
 292 versus initial condition based studies.

293 Another important factor is the realisation that the L_1 -norm is notoriously
 294 hard to optimise due to its non-smoothness near the origin. Intuitively, we can
 295 visualise the problem by considering the unconstrained optimisation problem of
 296 minimising the L_1 -norm of a scalar a . Using our gradient based approach, this
 297 amounts to stepping in the direction of steepest descent, which for our simple
 298 example is the sign of a . No matter how near or far we are to the optimal value
 299 of $a = 0$, this gradient will have the same magnitude. This means that we will
 300 continuously step over the optimal value, unless the step-size is perfect, leading
 301 to zig-zagging and ultimately causing the algorithm to converge rather slowly.
 302 To alleviate this behavior we replace the L_1 -norm with a smooth approximation,
 303 namely $l_{1,\delta}(\mathbf{q}) = h_\delta(\mathbf{q})/\delta$ where $h_\delta(\mathbf{q})$ is the pseudo-Huber norm (Bube & Langan
 304 1997; Bube & Nemeth 2007)

$$305 \quad h_\delta(\mathbf{q}) = \sum_j \delta^2 \left(\sqrt{1 + \frac{|q_j|^2}{\delta^2}} - 1 \right). \quad (2.14)$$

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

$$310 \quad \sigma_\delta^{(L_1)} = \frac{\|\mathbf{M}_q \mathcal{H} \mathbf{M}_f^{-1} \mathbf{f}_M\|_2}{l_{1,\delta}(\mathbf{T}(\mathbf{f}_M))}, \quad (2.15)$$

311 for decreasing values of δ . By using the optimal forcing obtained from an optimi-
 312 sation for the preceding one with a lower value of δ , we are able to more robustly
 313 achieve a converged optimisation for a sufficiently small δ such that our norm
 314 (2.14) is an appropriate approximation for the true L_1 -norm.

315 Before concluding this section, it is important to note that our choice of cost
 316 functional is not unique. Indeed, other cost functionals such as $\sigma^{(L_1)} = \sigma^{(L_2)} -$
 317 $\mu \|\mathbf{T}(\mathbf{f}_M)\|_1$ can also lead to sparse forcing modes for appropriate choices of μ .
 318 However, the fact that unit-norm forcings can lead to gains in energy many orders
 319 of magnitude larger than that of the forcing makes the choice of μ , which must
 320 balance the L_2 -based gain against the L_1 -based forcing, a difficult challenge.
 321 This is further complicated by the strong dependence of the gain on the forcing
 322 frequency, making a universally good way of choosing μ hard to determine. In
 323 our proposed cost functional there is no such parameter to choose, meaning that
 324 it can easily be applied to different frequencies and base-flows without change.
 325 Hence, we continue with it for the rest of the study.

326 *2.3. Resolvent wavemaker*

327 One concept that we use in our subsequent analysis is that of structural sensitivity
 328 and the wavemaker (Giannetti & Luchini 2007). The wavemaker has its origin
 329 in global stability analysis and provides a way to highlight regions in which flow
 330 field changes result in changes to global modes. Specifically, the wavemaker is the
 331 structural sensitivity to a localised flow feedback. To obtain the wavemaker, we
 332 consider the eigenvalue problem $\mathbf{L}\mathbf{x} = \lambda\mathbf{G}\mathbf{x}$. This eigenvalue problem could arise,
 333 for instance, as a global stability problem, in which case \mathbf{x} would be the global
 334 mode, with the corresponding eigenvalue λ giving its frequency and growth rate.
 335 It can be shown (see the review of Schmid & Brandt (2014) for example) that to
 336 first order, a perturbation to the eigenvalue $\delta\lambda$ for a perturbation in the matrix
 337 $\delta\mathbf{L}$ is given via

$$338 \quad \delta\lambda = \frac{\langle \mathbf{x}^\dagger, \delta\mathbf{L}\mathbf{x} \rangle}{\langle \mathbf{x}^\dagger, \mathbf{G}\mathbf{x} \rangle}, \quad (2.16)$$

339 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
 340 wavemaker is then obtained by specifying $\delta\mathbf{L} = \mathbf{I}$ and instead taking the element-
 341 wise, or Hadamard (\odot), product.

$$342 \quad \boldsymbol{\lambda} = \frac{\bar{\mathbf{x}}^\dagger \odot \mathbf{x}}{\langle \mathbf{x}^\dagger, \mathbf{G}\mathbf{x} \rangle}. \quad (2.17)$$

343 In this way, the wavemaker $\boldsymbol{\lambda}$ can then be thought of as a vector-field $\boldsymbol{\lambda}(x, y) =$
 344 $(\lambda_u, \lambda_v, \lambda_w)$ whose components represent what changes to the eigenvalue occur
 345 from localised feedback at each location and state-component in the flow field.

346 We quickly note that there are a few ways in which the wavemaker could be
 347 perceived. Whilst we have stayed within a discrete setting Giannetti & Luchini
 348 (2007) present the wavemaker in a continuous formulation. This gives the main

349 difference that their wavemaker is a scalar field which is defined pointwise via
 350 $\lambda(x, y) = \|\mathbf{x}^\dagger(x, y)\| \|\mathbf{x}(x, y)\|$. Hence, their wavemaker, by the Cauchy-Schwarz
 351 inequality, shows the maximum change to the eigenvalue that can be achieved via
 352 localised feedback at each spatial location. Conversely, the wavemaker presented
 353 by Schmid & Brandt (2014) is more easily related to ours via $\lambda(x, y) = \lambda_u +$
 354 $\lambda_v + \lambda_w$. In this manner, they obtain a complex-valued wavemaker whose real
 355 and imaginary parts show the individual changes to the real and imaginary parts
 356 of the eigenvalue. Additionally, the sign of these changes is retained, allowing
 357 for the direction of the eigenvalue perturbation to be determined. However, by
 358 keeping the values of the flow-field separate, our wavemaker definition is strongly
 359 related to that of Paladini *et al.* (2019), who introduce windowing matrices to
 360 allow for the selection of specific physical components in the resulting wavemaker.
 361 Whilst they use these matrices to isolate the contribution of the momentum to the
 362 wavemaker, we instead do this procedure for each separate velocity component.
 363 This means that for each spatial location, our wavemaker tells us how a specific
 364 eigenvalue will move for localised feedback restricted to each component of the
 365 state-vector.

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

$$371 \quad \delta\sigma = \sigma^2 \text{Real}(\langle \mathbf{f}, \delta \mathbf{L} \mathbf{q} \rangle_{\mathbf{w}_f}), \quad (2.18)$$

372 where \mathbf{L} stands for the linearised Navier–Stokes operator (Fosas de Pando *et al.*
 373 2014; Fosas de Pando & Schmid 2017; Qadri & Schmid 2017). Again, taking
 374 $\delta \mathbf{L} = \mathbf{I}$ and using the Hadamard product yields

$$375 \quad \sigma = \sigma^2 \text{Real}(\bar{\mathbf{f}} \odot \mathbf{W}_f \mathbf{q}). \quad (2.19)$$

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

380 An example of the wavemaker and resolvent wavemaker is shown for cylinder
 381 flow in figure 3. It is important to note that for the eigenvalue-based wavemaker
 382 the frequency is set by the eigenvalues. However, our definition of the resolvent
 383 wavemaker allows any frequency to be specified. Therefore, we concentrate on
 384 $St = 0.162$, which is the frequency at which the most unstable eigenvalue
 385 is found. We observe that these wavemakers have similar structures but with
 386 different gains. The fact that they have similar structures is not surprising, since
 387 the resolvent forcing and response modes are qualitatively similar to the direct
 388 and adjoint eigenvectors, respectively. However, the signs of the structures are
 389 often different. This indicates that a localised feedback affects the eigenvalue
 390 perturbation differently from the singular value perturbation, highlighting the
 391 importance of formulating a resolvent wavemaker in order to quantify the effect
 392 of localised feedback for resolvent analyses.




Figure 3: The eigenvalue-based wavemakers and resolvent wavemakers for cylinder flow at $Re = 100$ shown for illustration. Computations performed with the immersed boundary projection method *ibmos* (Fosas de Pando 2020).

393 3. Numerical setup

394 This section describes the numerical setup for our flow examples. In addition to
 395 the details given in this section, all Riemannian optimisations are managed using
 396 the python package *pyManopt* (Townsend *et al.* 2016), the python extension of
 397 the MATLAB package *Manopt* (Boumal *et al.* 2014). The optimisations are all
 398 conducted using the conjugate gradient algorithm.

399

3.1. Plane Poiseuille flow

400 Firstly, we consider plane Poiseuille flow to compare the differences in localisation
 401 strategies between initial-condition-based (transient growth) and driven (resol-
 402 vent) studies. The present setup follows that of Foures *et al.* (2013) in which
 403 a domain covers $(x, y) \in [0, 2\pi] \times [0, 2]$ at a Reynolds number of $Re = 4000$.
 404 No-slip boundary conditions are applied at $y = 0$ and 2 and periodic boundary
 405 conditions are applied at $x = 0$ and 2π . The base flow is analytically provided as
 406 $u = y(2 - y)$. We conduct the numerical simulations using the python package
 407 *ibmos* developed by Fosas de Pando (2020). This is an immersed boundary
 408 projection code based on the formulation of Taira & Colonius (2007) with specific
 409 formulation for optimisation and stability analyses. The package solves the non-
 410 linear incompressible Navier–Stokes equations and directly provides the linearised
 411 and adjoint codes necessary for conducting a resolvent analysis. For the plane
 412 Poiseuille examples, the matrix \mathbf{B} is chosen so that the forcing is only added to

413 the momentum equations. Similarly, the matrix \mathbf{C} is chosen so that only velocity
414 components constitute the output.

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

430 To this end, we consider a vector of the following form: $\mathbf{T}(\mathbf{f}_M) = \mathbf{M}(\mathbf{u} \odot \mathbf{u} +$
431 $\mathbf{v} \odot \mathbf{v})^{1/2}$, where \odot is the Hadamard product and the square root is taken compo-
432 nentwise. This vector has the same 2-norm as \mathbf{f}_M , but groups local contributions
433 of the forcing mode to the total energy together. Hence, the L_1 -norm of this
434 vector is small when the forcing is localised in space, but without penalising
435 among individual components of the velocity vector. It should be noted that
436 some additional care may be needed when designing this vector depending on
437 the specific numerical implementation. For example, as our immersed boundary
438 implementation uses a staggered mesh with the x -components of velocity lying
439 on the east and west faces of the cell whilst the y -components lie on the north
440 and south faces, we form the vector $\mathbf{T}(\mathbf{f}_M)$ on the cell centres by averaging
441 the kinetic energy contributions from the cell-edges. The weight matrices are
442 chosen to incorporate the grid spacing (see Taira & Colonius (2007) for more
443 information), so that the forcing and response are measured in terms of the
444 kinetic energy.

445

3.2. Flow past an aerofoil

446 Secondly, we also consider a spanwise-periodic turbulent flow over a canonical
447 aerofoil obtained from a large-eddy simulation (LES) with a Vremen sub-grid
448 scale model (Vremen 2004). The LES is conducted using the finite-volume solver
449 *CharLES* that solves the compressible Navier–Stokes equations with second-order
450 spatial and third-order temporal accuracies (Khalighi *et al.* 2011; Brès *et al.*
451 2017). The linearisation is performed within the same solver (Sun *et al.* 2017),
452 considering the time- and spanwise-averaged turbulent flow over the aerofoil as
453 the base flow.

454 The resolvent analysis is performed on a separate mesh from that used by
455 the LES. The mesh for the resolvent analysis has a two-dimensional rectangular
456 domain with the extent of $x/L_c \in [-15, 16]$ and $y/L_c \in [-6, 5]$, comprising ap-
457 proximately 0.11 million cells and giving the resulting discretised linear operator a
458 dimension of 540840×540840 . Compared to the LES mesh, the mesh for resolvent
459 analysis is coarser over the aerofoil and in the wake, but is much finer in the
460 upstream region of the aerofoil in order to resolve the forcing mode structures.
461 The convergence of resolvent norm with respect to the domain extent and grid

462 resolution has been reported in detail in Yeh & Taira (2019). At the far-field
 463 boundary and over the aerofoil, Dirichlet conditions are specified for the density
 464 and velocities and Neumann conditions are prescribed for the pressure in \mathbf{q} . At
 465 the outlet boundary, Neumann conditions are provided for all flow variables. The
 466 base-flow is two-dimensional, however, in contrast to the plane Poiseuille case,
 467 and we allow the perturbations to be three-dimensional by adopting a bi-global
 468 setting that decomposes \mathbf{q} into spanwise Fourier modes with the wavenumber β .

469 Even though the linear operator is sparse, its large dimension requires special
 470 care. To efficiently deal with this operator, the python bindings for PETSc (Balay
 471 *et al.* 2021a,b, 1997), `petsc4py` (Dalcin *et al.* 2011) are used. This enables us
 472 to carry out the required linear algebra manipulations in parallel whilst keeping
 473 our code within the python environment. Specifically, PETSc is used together
 474 with the external library MUMPS (Amestoy *et al.* 2001, 2019) in order to provide
 475 the LU decomposition of the resolvent operator, and hence evaluate the action
 476 of the resolvent operator (and its adjoint) on a vector. To compare our sparse
 477 method with a traditional resolvent analysis, the SVD of the resolvent operator
 478 is found using a Lanczos SVD solver provided by the python bindings for SLEPc
 479 (Hernandez & Vidal 2005), `slepc4py` (Dalcin *et al.* 2011).

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

$$486 \quad \|\mathbf{q}\|_E^2 = \int_{\Omega} \left(\frac{RT_0}{\rho_0} |\rho'|^2 + \rho_0 \|\mathbf{u}'\|^2 + \frac{R\rho_0}{(\gamma-1)T_0} |T'|^2 \right) dV, \quad (3.1)$$

487 which represents the energy contained in a perturbation in the absence of
 488 compression work. In defining the Chu-norm, we have used a dash ' to denote
 489 quantities derived from our state vector \mathbf{q} . Similarly a subscript 0 is used
 490 to denote quantities derived from the base-flow used for linearisation. This
 491 integral is discretised to $\|\mathbf{q}\|_E^2 = \mathbf{q}^H \mathbf{W}_E \mathbf{q}$ with \mathbf{W}_E as a positive definite
 492 weight matrix. Taking the Cholesky decomposition $\mathbf{W}_E = \mathbf{M}^H \mathbf{M}$ gives the
 493 matrices needed for the resolvent description ($\mathbf{M}_q = \mathbf{M}_f = \mathbf{M}$). Hence, to
 494 keep momentum grouped in our sparsification, we split the components of
 495 the norm matrix \mathbf{M} to form the state $\mathbf{T}(\mathbf{f}_M) = (\mathbf{M}_\rho \rho', \mathbf{M}_{KE} \boldsymbol{\kappa}', \mathbf{M}_T T')$ where
 496 $\boldsymbol{\kappa}' = \sqrt{(\rho\mathbf{u})' \odot (\rho\mathbf{u})' + (\rho\mathbf{v})' \odot (\rho\mathbf{v})' + (\rho\mathbf{w})' \odot (\rho\mathbf{w})'}$. Note that in defining $\boldsymbol{\kappa}'$,
 497 we have used the notation $(\rho\mathbf{u})' = \rho_0 \mathbf{u}' + \rho' \mathbf{u}_0$ for the streamwise linearised
 498 momentum component, with similar definitions for the spanwise $(\rho\mathbf{v})'$ and
 499 transverse $(\rho\mathbf{w})'$ linearised momentum components. This vector has the same
 500 L_2 -norm as our full state vector. However, when taking the L_1 -norm the kinetic
 501 energy is now grouped, ensuring that all velocity components are treated equally.
 502 It should be pointed out that now we have two additional components in the
 503 norm, specifically ρ' and T' . By not including these in the same component of
 504 the vector for the L_1 -norm, they are treated separately by the sparsification. In
 505 essence, this means that the sparse resolvent does not only sparsify the spatial
 506 structure of any forcing but also sparsifies the actuation mechanism by choosing
 507 between a velocity-based, density-based or temperature-based forcing.




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

508 4. Results

509

4.1. Plane Poiseuille flow

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

$$516 \quad \mathbf{f}(x, y) = \mathbf{f}_\alpha(y)e^{i\alpha x} \quad \text{and} \quad \mathbf{u}(x, y) = \mathbf{u}_\alpha(y)e^{i\alpha x}. \quad (4.1)$$

517 However, as we are using a global (2D) analysis, the wavenumbers that our forcing
 518 and response can consist of are set by the aspect ratio of the domain. Taking
 519 $x \in [0, 2\pi]$ with periodic boundary conditions requires our wavenumbers to be
 520 integer, i.e., $\alpha \in \mathbb{N}$. Another artifact of using a 2D code for a parallel-flow is that
 521 the results do not change if the forcing and response modes are translated along
 522 the x -axis.

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




Figure 5: The real parts of the forcing and response mode 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.




Figure 6: The real parts of the forcing and response mode 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.

539 is also evident from the forcing mode being situated more centrally in the y -
 540 direction where the base-flow has a higher velocity.

541 We now turn our attention to the resolvent analysis results from the sparse
 542 optimisation procedure. The sparse forcing mode obtained for $\omega = 0.278$ is
 543 shown in figure 7. Firstly, it is clear from this figure that the forcing mode is
 544 more sparse than the full resolvent analysis. Indeed, instead of a series of slanted
 545 structures angled against the shear we now have thin stripes parallel to the walls.
 546 Interestingly, even though our vector $\mathbf{T}(\mathbf{f}_M)$ was carefully chosen not to sparsify
 547 the separate velocity components at a given spatial location, the sparse forcing
 548 mode consists mainly of a u -component, indicating that this forcing is primarily
 549 in the direction of the wall.

550 A striking feature of the sparse forcing mode is that it has maintained its

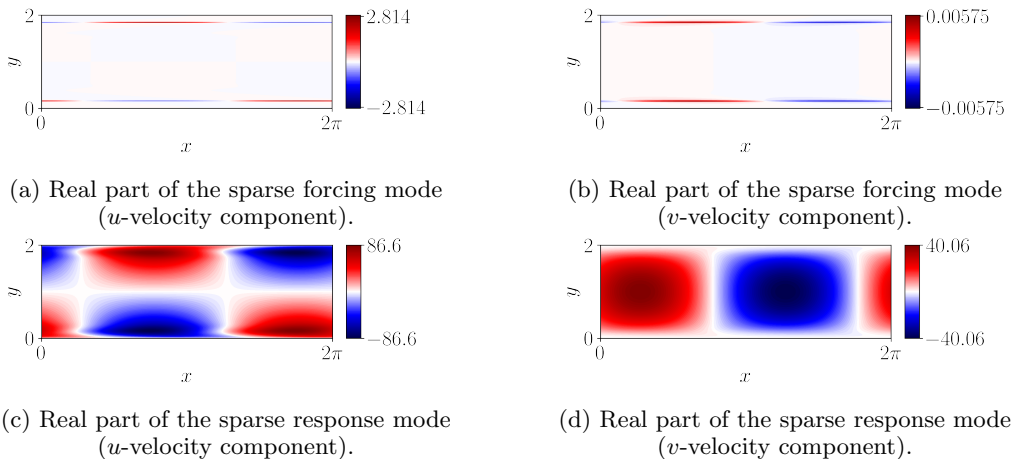


Figure 7: The real parts of the forcing and response mode obtained by sparsification at $\omega = 0.278$. The forcing is unit-norm whereas the response mode has norm equal to the gain.

551 $\alpha = 1$ structure, i.e. it is still 2π -periodic in the x -direction. This is particularly
 552 enlightening since the strip-structure is not as sparse as the forcing mode could
 553 be, which would consist of just one element of the kinetic energy vector being
 554 filled, i.e. a single spatial location forcing. Therefore, the fact that the sparse
 555 procedure has chosen a less sparse structure indicates that forcing with this spatial
 556 wavenumber is crucial in achieving a high gain at this frequency. The location of
 557 these stripes can be hypothesised to be intrinsically linked to the $\alpha = 1$ structure
 558 using the concept of critical layers. A critical layer occurs at y^* where the base-
 559 flow velocity $U(y^*)$ is equal to the phase velocity k of a disturbance, and is central
 560 in causing instability in plane Poiseuille flow. Using the phase velocity k for a
 561 disturbance at $\omega = 0.278$ and $\alpha = 1$ implies a critical layer at $y^* \approx 0.150$, which
 562 is close to the y -location of the stripes which occur at $y = 0.155$. Hence, the
 563 sparse forcing mechanism can be summarised as forcing along (or just above)
 564 and parallel to the critical layer, with the v -velocity component, which does not
 565 contribute to this critical layer, being negligible.

566 Further evidence for the importance of $\alpha = 1$ forcing is shown in the response
 567 modes, which are also displayed in figure 7. The figure shows that the response
 568 modes stemming from the sparse forcing mode have the same structure as those
 569 from the full resolvent. As well as reinforcing that the $\alpha = 1$ forcing is critical for
 570 providing optimal amplification at this frequency, this observation also highlights
 571 the low-rank nature of the resolvent at this frequency. Even though the forcing
 572 shape is qualitatively different in the sparse case, the shape of the response is
 573 identical, disregarding arbitrary phase shifts, albeit with a lower magnitude. This
 574 agrees with previous observations, such as that of Rosenberg *et al.* (2019) where
 575 it is shown that when there is a large separation in singular values the shape of
 576 the forcing is less critical in exciting the dominant response. The lower magnitude
 577 is to be expected since our sparse forcing mode sacrifices some amount of energy
 578 to achieve a more localised spatial structure.

579 To provide a comparison between the results at different wavenumbers, we also
 580 carry out the sparse optimisation procedure at $\omega = 1.14$. Figure 8 shows the
 581 results, which differ quite significantly from the case of $\omega = 0.278$. In this case,




Figure 8: The real parts of the forcing and response mode 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 .

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

600

4.2. Flow past an aerofoil

601 Secondly, we consider flow past a NACA 0012 aerofoil at an angle of attack
 602 of 9° , a chord-based Reynolds number of $Re = 23000$ and a free stream Mach
 603 number of $M = 0.3$ (see section 3.2 for the numerical details). In contrast to
 604 the plane Poiseuille example, this flow is unsteady and turbulent. Therefore,
 605 the mean-flow is used for linearisation. This time-averaged base-flow is shown
 606 in figure 9. Similarly to the work of Yeh & Taira (2019), and Ribeiro *et al.*
 607 (2020) who considered a resolvent analysis with the same base flow, we consider
 608 the resolvent modes at spanwise wavenumbers $\beta = 0$ and $20/\pi$. As the linear
 609 operator is unstable, a discounting parameter of $\alpha = 0.63$ is used (Jovanović




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 $St \approx 5.22$ for $\beta = 0$ and $St \approx 5.90$ for $\beta = 20\pi$ (both highlighted with a black +).

610 2004). The gain-frequency relationships are shown in figure 10. Similarly to the
 611 previous examples, we focus our subsequent analysis on the frequencies at which
 612 the peak gain is obtained.

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

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

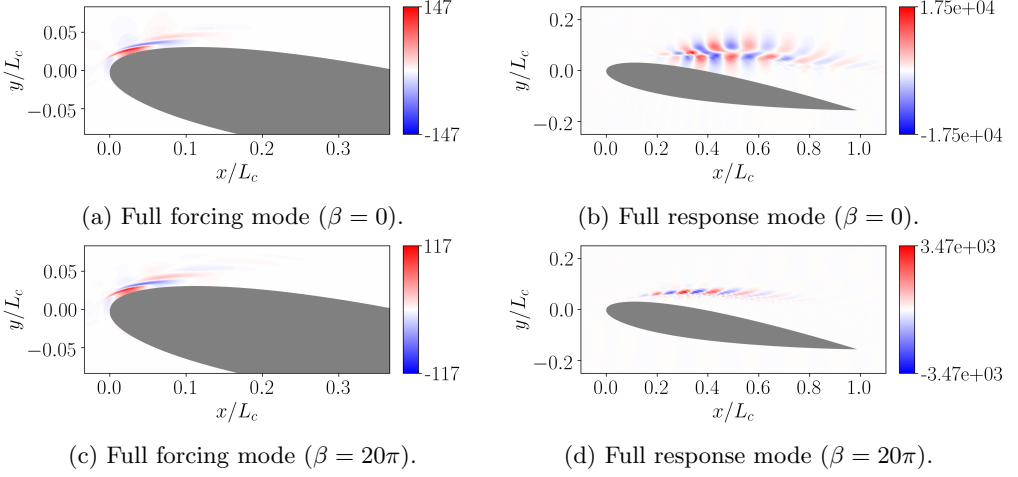


Figure 11: The full resolvent modes for aerofoil flow at $St = 5.22$ for $\beta = 0$ and $St = 5.90$ for $\beta = 20\pi$. The linearised component of the streamwise momentum is shown.

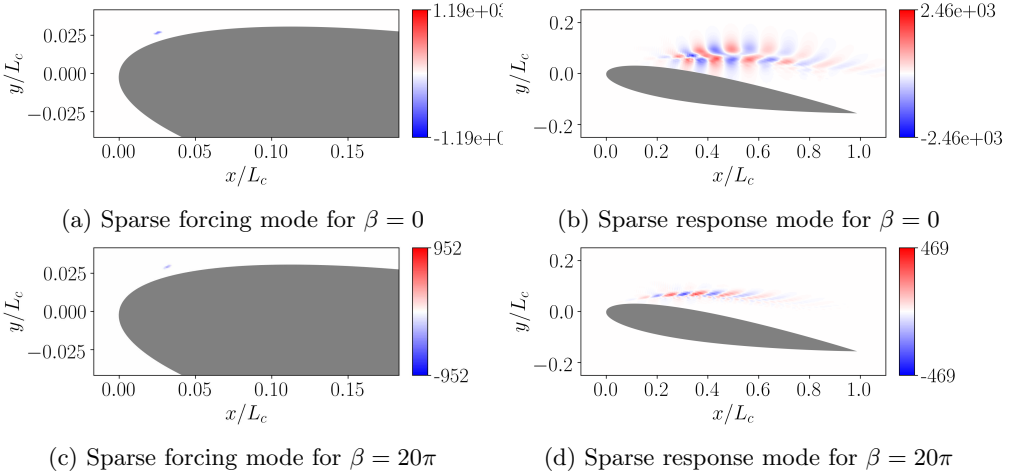


Figure 12: The sparse resolvent modes for aerofoil flow at $St = 5.22$ for $\beta = 0$ and $St = 5.90$ for $\beta = 20\pi$. The linearised component of the streamwise momentum is shown.

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

637 To provide additional insight into the chosen spatial location, we now examine
 638 the resolvent wavemakers, which are shown in figure 13. For both values of β ,
 639 the wavemakers display a large positive region in the mean-flow shear layer.
 640 This is not unexpected, since regions of shear translate to regions of high non-




Figure 13: The resolvent wavenakers 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.




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.

641 normality in the linearised Navier–Stokes operator, which underpins sensitive
 642 areas for forcing. In fact, in the full-SVD forcing modes, we directly see this,
 643 as the forcing modes in both cases are primarily located in this shear region.
 644 The sparse forcing locations are also situated in this region and are located near
 645 the maximum value of the full-SVD forcing mode. Whilst this may show that
 646 choosing the largest value of the forcing mode is a good candidate for the sparse
 647 forcing mode, we emphasise that the optimisation procedure is not biased by
 648 any knowledge of the full forcing mode, and all physical mechanisms and spatial
 649 locations are weighted equally.

650 Whilst forcing in the shear region may provide the optimal response, it is rather
 651 impractical for flow actuation purposes. Therefore, we conclude this section by
 652 using the windowing matrix \mathbf{B} to conduct our analysis on the surface of the
 653 aerofoil. Figure 14 shows the forcing distribution along the suction side of the
 654 aerofoil as a function of the distance along the chord x_c . In the full resolvent
 655 analysis, most of the forcing is concentrated near the leading edge of the aerofoil,
 656 agreeing with the non-windowed case. In both spanwise wavenumber cases, the
 657 sparse mode is once again a single-point momentum-based forcing and is located
 658 at the maximum value for the kinetic energy of the full forcing mode. This
 659 provides the optimal compromise between forcing with $(\rho u)'$ at its maximum
 660 value and $(\rho v)'$ at its maximum, which in the full case is located to the left of
 661 $(\rho u)'$. Even though there is a $(\rho E)'$ -component, this is simply a consequence of the
 662 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
 663 is no thermodynamic contribution to the linearised total-energy. The importance

664 of grouping momentum together into one coherent strategy is evident from the
 665 figure, as the directional information of the actuation is crucial in both the full
 666 and sparse resolvent analyses to achieve the optimal gain. This information would
 667 otherwise be lost. Moreover, by grouping the momentum together we strike a
 668 compromise between choosing a forcing that is optimal for each isolated velocity
 669 component.

670 5. Conclusion

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

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

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

703 Overall, the sparse optimisation procedure provides an unbiased optimal spar-
 704 sification of the flow and is able to adapt to the different forcing strategies
 705 available at different frequencies. Although the aerofoil results show that choosing
 706 the maxima of the SVD is a good candidate for a sparse forcing vector, the
 707 plane Poiseuille example shows that both single-point and multi-point forcing
 708 modes can be found depending on the low-rank nature and physical mechanisms
 709 furnished by the resolvent. Based on our results, it can be postulated that in
 710 more complex systems, such as those stemming from aeroacoustic or combustion
 711 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 further note that recent efforts have been made to extend resolvent analysis to both periodic flows (Padovan *et al.* 2020) and also to the non-linear regime (Rigas *et al.* 2021). As the techniques 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.

Funding. KT and PJS thank the support from the Army Research Office (grant: W911NF-21-1-0060, program officer: Matthew J. Munson). KT also acknowledges the support from the Office of Naval Research (grant: N0014-19-1-2460, program officer: David R. Gonzalez).

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

Author ORCID. C. S. Skene, <https://orcid.org/0000-0003-0994-2013>;

C.-A. Yeh, <https://orcid.org/0000-0003-0426-8381>;

P. J. Schmid, <https://orcid.org/0000-0002-2257-8490>;


K. Taira, <https://orcid.org/0000-0002-3762-8075>;

Appendix A. 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 section 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 15a and 15c 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 15b and 15d, 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.


Also shown in figures 16a and 16b is the norm of the gradient provided by the optimisation procedure as a function of the number of iterations. For the case of $\omega = 0.278$, 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




(a) L_1 norm and gains with delta
($\omega = 0.278$)



(b) Convergence of the pseudo-Huber norm to the true L_1 norm ($\omega = 0.278$)




(c) L_1 norm and gains with delta
($\omega = 1.14$)




(d) Convergence of the pseudo-Huber norm to the true L_1 norm ($\omega = 1.14$)

Figure 15: Dependence of the results on δ for the plane Poiseuille flow examples



(a) Convergence behaviour of the conjugate gradient algorithm for plane Poiseuille flow at $\omega = 0.278$ and $\delta = 5^{-7}$.



(b) Convergence behaviour of the conjugate gradient algorithm for plane Poiseuille flow at $\omega = 1.14$ and $\delta = 5^{-7}$.

Figure 16

760 the difficulty of the gradient-based optimisation. This is especially highlighted
 761 by figure 16b, where the optimisation terminates earlier due to a stagnation of
 762 the cost functional. This showcases the importance of using both a relaxation
 763 parameter δ as well as an optimisation procedure such as the conjugate gradient
 764 algorithm in order to achieve converged results.

- 765 ABSIL, P.-A., MAHONY, R. & SEPULCHRE, R. 2007 *Optimization Algorithms on Matrix*
766 *Manifolds*. USA: Princeton University Press.
- 767 AMESTOY, P.R., BUTTARI, A., L'EXCELLENT, J.-Y. & MARY, T. 2019 Performance and
768 Scalability of the Block Low-Rank Multifrontal Factorization on Multicore Architectures.
769 *ACM Transactions on Mathematical Software* **45**, 2:1–2:26.
- 770 AMESTOY, P.R., DUFF, I. S., KOSTER, J. & L'EXCELLENT, J.-Y. 2001 A fully asynchronous
771 multifrontal solver using distributed dynamic scheduling. *SIAM Journal on Matrix*
772 *Analysis and Applications* **23** (1), 15–41.
- 773 BALAY, S., ABHYANKAR, S., ADAMS, M. F., BROWN, J., BRUNE, P., BUSCHELMAN, K.,
774 DALCIN, L., DENER, A., ELJKHOUT, V., GROPP, W. D., KAUSHIK, D., KNEPLEY, M. G.,
775 MAY, D. A., MCINNES, L. C., MILLS, R. T., MUNSON, T., RUPP, K., SANAN, P., SMITH,
776 B. F., ZAMPINI, S., ZHANG, H. & ZHANG, H. 2021a PETSc users manual. *Tech. Rep.*
777 ANL-95/11 - Revision 3.15. Argonne National Laboratory.
- 778 BALAY, S., ABHYANKAR, S., ADAMS, M. F., BROWN, J., BRUNE, P., BUSCHELMAN, K.,
779 DALCIN, L., DENER, A., ELJKHOUT, V., GROPP, W. D., KAUSHIK, D., KNEPLEY,
780 M. G., MAY, D. A., MCINNES, L. C., MILLS, R. T., MUNSON, T., RUPP, K., SANAN,
781 P., SMITH, B. F., ZAMPINI, S., ZHANG, H. & ZHANG, H. 2021b PETSc Web page.
782 <https://petsc.org/>.
- 783 BALAY, S., GROPP, W. D., MCINNES, L. C. & SMITH, B. F. 1997 Efficient management of
784 parallelism in object oriented numerical software libraries. In *Modern Software Tools in*
785 *Scientific Computing* (ed. E. Arge, A. M. Bruaset & H. P. Langtangen), pp. 163–202.
786 Birkhauser Press.
- 787 BOUMAL, N. & ABSIL, P.-A. 2015 Low-rank matrix completion via preconditioned optimization
788 on the Grassmann manifold. *Linear Algebra and its Applications* **475**, 200–239.
- 789 BOUMAL, N., MISHRA, B., ABSIL, P.-A. & SEPULCHRE, R. 2014 Manopt, a Matlab toolbox for
790 optimization on manifolds. *Journal of Machine Learning Research* **15**, 1455–1459.
- 791 BRÈS, G. A., HAM, F. E., NICHOLS, J. W. & LELE, S. K. 2017 Unstructured large-eddy
792 simulations of supersonic jets. *AIAA Journal* **55** (4), 1164–1184.
- 793 BUBE, K. P. & LANGAN, R. T. 1997 Hybrid ℓ_1 / ℓ_2 minimization with applications to
794 tomography. *GEOPHYSICS* **62** (4), 1183–1195.
- 795 BUBE, K. P. & NEMETH, T. 2007 Fast line searches for the robust solution of linear systems in
796 the hybrid ℓ_1 / ℓ_2 and Huber norms. *GEOPHYSICS* **72** (2), A13–A17.
- 797 CATTAFESTA, L. N. & SHEPLAK, M. 2011 Actuators for active flow control. *Annu. Rev. Fluid*
798 *Mech.* **43**, 247–272.
- 799 CHU, B. T. 1965 On the energy transfer to small disturbances in fluid flow (part I). *Acta*
800 *Mechanica* **1** (3), 215–234.
- 801 DALCIN, L. D., PAZ, R. R., KLER, P. A. & COSIMO, A. 2011 Parallel distributed computing
802 using python. *Advances in Water Resources* **34** (9), 1124 – 1139, New Computational
803 Methods and Software Tools.
- 804 FARRELL, B. F. & IOANNOU, P. J. 1993 Stochastic forcing of the linearized Navier–Stokes
805 equations. *Physics of Fluids A: Fluid Dynamics* **5** (11), 2600–2609.
- 806 FOSAS DE PANDO, M. 2020 IBMOS: Immersed boundary method optimization and stability.
807 <https://doi.org/10.5281/zenodo.3757783>.
- 808 FOSAS DE PANDO, M. & SCHMID, P. J. 2017 Optimal frequency-response sensitivity of
809 compressible flow over roughness elements. *Journal of Turbulence* **18** (4), 338–351.
- 810 FOSAS DE PANDO, M., SCHMID, P. J. & LELE, S. K. 2014 Parametric sensitivity for large-scale
811 aeroacoustic flows. In *Proceedings of the 2014 Summer Program*, pp. 365–374. Center for
812 Turbulence Research, Stanford University.
- 813 FOURES, D. P. G., CAULFIELD, C. P. & SCHMID, P. J. 2013 Localization of flow structures
814 using ∞ -norm optimization. *Journal of Fluid Mechanics* **729**, 672–701.
- 815 GIANNETTI, F. & LUCHINI, P. 2007 Structural sensitivity of the first instability of the cylinder
816 wake. *Journal of Fluid Mechanics* **581**, 167–197.
- 817 HERNANDEZ, V. AND ROMAN, J. E. & VIDAL, V. 2005 Slepnc: A scalable and flexible toolkit for
818 the solution of eigenvalue problems. *ACM Trans. Math. Softw.* **31** (3), 351–362.
- 819 HUANG, W., ABSIL, P.-A. & GALLIVAN, K. A. 2016 A Riemannian BFGS method for
820 nonconvex optimization problems. In *Numerical Mathematics and Advanced Applications*
821 *ENUMATH 2015* (ed. B. Karasözen, M. Manguoğlu, M. Tezer-Sezgin, S. Göktepe &

- 822 Ö. Uğur), *Lecture Notes in Computational Science and Engineering* 112, pp. 627–634.
823 Cham: Springer International Publishing.
- 824 JEUN, J., NICHOLS, J. W. & JOVANOVIĆ, M. R. 2016 Input-output analysis of high-
825 speed axisymmetric isothermal jet noise. *Physics of Fluids* **28** (4), 047101, arXiv:
826 <https://doi.org/10.1063/1.4946886>.
- 827 JOVANOVIĆ, M. R. 2004 Modeling, analysis, and control of spatially distributed systems. Ph. D.
828 thesis, University of California, Santa Barbara.
- 829 JOVANOVIĆ, M. R. & BAMIEH, B. 2005 Componentwise energy amplification in channel flows.
830 *Journal of Fluid Mechanics* **534**, 145–183.
- 831 KHALIGHI, Y., NICHOLS, J. W., HAM, F., LELE, S. K. & MOIN, P. 2011 Unstructured large
832 eddy simulation for prediction of noise issued from turbulent jets in various configurations.
833 AIAA Paper 2011-2886.
- 834 LIU, Q., SUN, Y., YEH, C.-A., UKEILEY, L. S., CATTAFESTA, L. N. & TAIRA, K. 2021 Unsteady
835 control of supersonic turbulent cavity flow based on resolvent analysis. *Journal of Fluid*
836 *Mechanics* **925**, A5.
- 837 LUHAR, M., SHARMA, A. S. & MCKEON, B. J. 2014 Opposition control within the resolvent
838 analysis framework. *Journal of Fluid Mechanics* **749**, 597–626.
- 839 MCKEON, B. J. & SHARMA, A. S. 2010 A critical-layer framework for turbulent pipe flow.
840 *Journal of Fluid Mechanics* **658**, 336–382.
- 841 PADOVAN, A., OTTO, S. E. & ROWLEY, C. W. 2020 Analysis of amplification mechanisms
842 and cross-frequency interactions in nonlinear flows via the harmonic resolvent. *Journal*
843 *of Fluid Mechanics* **900**, A14.
- 844 PALADINI, E., BENEDDINE, S., DANDOIS, J., SIPP, D. & ROBINET, J.-C. 2019 Transonic buffet
845 instability: From two-dimensional airfoils to three-dimensional swept wings. *Phys. Rev.*
846 *Fluids* **4**, 103906.
- 847 PRINGLE, C. C. T., WILLIS, A. P. & KERSWELL, R. R. 2012 Minimal seeds for shear flow
848 turbulence: using nonlinear transient growth to touch the edge of chaos. *Journal of Fluid*
849 *Mechanics* **702**, 415–443.
- 850 QADRI, U. A. & SCHMID, P. J. 2017 Frequency selection mechanisms in the flow of a laminar
851 boundary layer over a shallow cavity. *Phys. Rev. Fluids* **2**, 013902.
- 852 RIBEIRO, J. H. M., YEH, C.-A. & TAIRA, K. 2020 Randomized resolvent analysis. *Phys. Rev.*
853 *Fluids* **5**, 033902.
- 854 RIGAS, G., SIPP, D. & COLONIUS, T. 2021 Nonlinear input/output analysis: application to
855 boundary layer transition. *Journal of Fluid Mechanics* **911**, A15.
- 856 ROSENBERG, K., SYMON, S. & MCKEON, B. J. 2019 Role of parasitic modes in nonlinear
857 closure via the resolvent feedback loop. *Phys. Rev. Fluids* **4**, 052601.
- 858 SCHMID, P. J. & BRANDT, L. 2014 Analysis of Fluid Systems: Stability, Receptivity, Sensitivity:
859 Lecture notes from the FLOW-NORDITA Summer School on Advanced Instability
860 Methods for Complex Flows, Stockholm, Sweden, 2013. *Applied Mechanics Reviews*
861 **66** (2), 024803.
- 862 SCHMID, P. J. & HENNINGSON, D. S. 2001 *Stability and Transition in Shear Flows, Applied*
863 *Mathematical Sciences*, vol. 142. New York, NY: Springer-Verlag.
- 864 SUN, Y., TAIRA, K., CATTAFESTA, L. N. & UKEILEY, L. S. 2017 Biglobal instabilities of
865 compressible open-cavity flows. *Journal of Fluid Mechanics* **826**, 270–301.
- 866 TAIRA, K., BRUNTON, S. L., DAWSON, S. T. M., ROWLEY, C. W., COLONIUS, T., MCKEON,
867 B. J., SCHMIDT, O. T., GORDEYEV, S., THEOFILIS, V. & UKEILEY, L. S. 2017 Modal
868 analysis of fluid flows: An overview. *AIAA J.* **55** (12), 4013–4041.
- 869 TAIRA, K. & COLONIUS, T. 2007 The immersed boundary method: A projection approach.
870 *Journal of Computational Physics* **225** (2), 2118 – 2137.
- 871 TAIRA, K., HEMATI, M. S., BRUNTON, S. L., SUN, Y., DURAISAMY, K., BAGHERI, S., DAWSON,
872 S. & YEH, C.-A. 2020 Modal analysis of fluid flows: Applications and outlook. *AIAA J.*
873 **58** (3), 998–1022.
- 874 TOEDTLI, S. S., LUHAR, M. & MCKEON, B. J. 2019 Predicting the response of turbulent
875 channel flow to varying-phase opposition control: Resolvent analysis as a tool for flow
876 control design. *Phys. Rev. Fluids* **4**, 073905.
- 877 TOWNSEND, J., KOEP, N. & WEICHWALD, S. 2016 Pymanopt: A python toolbox for

- 878 optimization on manifolds using automatic differentiation. *Journal of Machine Learning*
879 *Research* **17** (137), 1–5.
- 880 TREFETHEN, L. N., TREFETHEN, A. E., REDDY, S. C. & DRISCOLL, T. A. 1993 Hydrodynamic
881 stability without eigenvalues. *Science* **261** (5121), 578–584.
- 882 VREMAN, A. W. 2004 An eddy-viscosity subgrid-scale model for turbulent shear flow: Algebraic
883 theory and applications. *Physics of Fluids* **16** (10), 3670–3681.
- 884 YEH, C.-A., BENTON, S. I., TAIRA, K. & GARMANN, D. J. 2020 Resolvent analysis of an airfoil
885 laminar separation bubble at $Re = 500000$. *Phys. Rev. Fluids* **5** (8), 1–24.
- 886 YEH, C.-A. & TAIRA, K. 2019 Resolvent-analysis-based design of airfoil separation control.
887 *Journal of Fluid Mechanics* **867**, 572–610.