



Beyond optimal disturbances: a statistical framework for transient growth

Peter Frame^{1,†} and Aaron Towne¹

¹Department of Mechanical Engineering, University of Michigan, Ann Arbor, MI, USA

(Received 13 March 2023; revised 26 September 2023; accepted 22 January 2024)

The theory of transient growth describes how linear mechanisms can cause temporary amplification of disturbances even when the linearized system is asymptotically stable as defined by its eigenvalues. This growth is traditionally quantified by finding the initial disturbance that generates the maximum response at the peak time of its evolution. However, this can vastly overstate the growth of a real disturbance. In this paper, we introduce a statistical perspective on transient growth that models statistics of the energy amplification of the disturbances. We derive a formula for the mean energy amplification and spatial correlation of the growing disturbance in terms of the spatial correlation of the initial disturbance. The eigendecomposition of the correlation provides the most prevalent structures, which are the statistical analogue of the standard left singular vectors of the evolution operator. We also derive accurate confidence bounds on the growth by approximating the probability density function of the energy. Applying our analysis to Poiseuille flow yields a number of observations. First, the mean energy amplification is often drastically smaller than the maximum. In these cases, it is exceedingly unlikely to achieve near-optimal growth due to the exponential behaviour observed in the probability density function. Second, the characteristic length scale of the initial disturbances has a significant impact on the expected growth, with large-scale initial disturbances growing orders of magnitude more than small-scale ones. Finally, while the optimal growth scales quadratically with Reynolds number, the mean energy amplification scales only linearly for certain reasonable choices of the initial correlation.

Key words: shear-flow instability, transition to turbulence

1. Introduction

A natural approach for analysing the stability of a steady fluid flow is to linearize and calculate the eigenvalues of the linearized Navier–Stokes operator. Underlying this

† Email address for correspondence: pframe@umich.edu

analysis is the assumption that there will be disturbances to the steady flow, and although their magnitude is difficult to know *a priori*, it is likely a value small enough that nonlinear mechanisms are not relevant. This approach, known as modal stability theory, is agnostic to the shape of any particular disturbance – if there is a positive eigenvalue, any disturbance arising in a physical scenario will grow, otherwise, any disturbance will decay asymptotically. However, the modal approach predicts stability when experiments tell us otherwise. Famously, Reynolds found that, at high velocities, pipe flow transitions to turbulence (Reynolds 1883). Efforts to ground this instability in modal theory floundered: pipe flow has all stable eigenvalues. The same is true for Couette flow as well as plane Poiseuille flow at low Reynolds numbers; these flows have only stable eigenvalues, but are observed to transition (Tillmark & Alfredsson 1992).

The key to their instability can be, in fact, a linear mechanism (Schmid 2007). Perhaps counterintuitively, a linearized Navier–Stokes operator with all stable eigenvalues can lead to short-term growth in the magnitude of disturbances before they decay at the rate prescribed by the least stable eigenvalue. This transient growth is possible only when the linearized Navier–Stokes operator is non-normal, i.e. its eigenvectors are not orthogonal in the energy norm. This permits one eigenvector to initially subtract from another, but this cancellation can cease if one eigenvector vanishes faster than the other, leading to growth. The magnitude of this growth can be remarkable – often more than one-thousand-fold at its peak (Trefethen *et al.* 1993). While the initial disturbances are assumed to be too small for nonlinear effects to be important, when they are amplified by three orders of magnitude, the assumption of linearity may no longer hold, and nonlinearities may bring the flow away from the laminar steady state. Rather than leading directly to transition, the nonlinearities activated by the amplified disturbance might bring the flow to a new state. Instability in this state, known as secondary instability, is more likely than the primary growth to lead to turbulence (Schmid & Henningson 2001).

The metric reported in the literature to quantify transient growth is the ratio of kinetic energy of the maximally amplified disturbance to its initial kinetic energy. This metric is usually referred to as $G(t)$, although in this paper we call it $G^{opt}(t)$ to distinguish it from suboptimal and mean growths. Significant effort has been devoted to studying $G^{opt}(t)$ both analytically and numerically. In channel flow, it can be shown to have quadratic dependence on the Reynolds number Re when the product of the streamwise wavenumber and Reynolds number is small, $\alpha Re \ll 1$ (Gustavsson 1991). Under the same conditions, the time at which the maximum occurs increases linearly with Re . Indeed, numerical experiments show that there is quadratic scaling in the optimal growth and linear scaling in the optimal time for plane Poiseuille (Trefethen *et al.* 1993), Couette (Trefethen *et al.* 1993), Blasius boundary layer (Butler & Farrell 1992; Hanifi, Schmid & Henningson 1996) and pipe (Schmid & Henningson 1994) flows. In all of these cases, the optimal streamwise wavenumber α is zero or very small, and the optimal spanwise wavenumber β is order unity (Schmid & Henningson 2001).

Minimal seed theory (Pringle, Willis & Kerswell 2012; Cherubini, De Palma & Robinet 2015; Kerswell 2018) provides a nonlinear analogue of optimal transient growth analysis. At each initial energy, it identifies the disturbance that achieves the greatest growth when evolved according to the full nonlinear Navier–Stokes equations. When the initial energy is just large enough that the optimal disturbance leads to sustained turbulence (when evolved for a sufficiently long period), the disturbance is called the minimal seed – the smallest disturbance leading to transition. Although minimal seeds are initially amplified by linear mechanisms (Pringle *et al.* 2012), they can differ substantially in shape from the optimal disturbances in linear transient growth (Pringle & Kerswell 2010). This gives a lower bound for the energy level that disturbances must achieve to spark transition.

In either the linear or nonlinear context, considering optimal disturbances gives an upper bound on the growth experienced by disturbances that are not influenced by further forcing, but we propose that a more complete picture of the possible growth is needed. In linear transient growth, only the optimal initial disturbance experiences G^{opt} growth. Indeed, if the initial disturbance were one of the eigenvectors of the linearized Navier–Stokes operator, it would decay monotonically. Of course, real disturbances to the flow will not exactly coincide with the optimally amplified disturbance, so in order to quantify their growth, one needs to explore the space of suboptimal disturbances. Is most of this space inhabited by disturbances that decay or by ones that grow? Is the growth of real disturbances of the order of G^{opt} , on average? What is the probability that a random disturbance will come close to G^{opt} ?

Motivated by these questions, we investigate transient growth from a statistical perspective in this paper. A statistical view serves both to model the uncertainty and variation in the spatial form of initial disturbances and to fully explore the high-dimensional space that these disturbances occupy. We derive an equation for the mean energy of the amplified random disturbances, and dividing this by the mean initial energy gives a metric for the mean energy amplification, which we term G^{mean} . This depends on the statistics of the incoming disturbances, and the formula we report for G^{mean} involves the correlation matrix of the initial disturbances. The correlation matrix of initial disturbances is distinct from the correlation matrix one measures in an experiment because the latter combines information about the initial disturbances and their evolution under the linear dynamics. The correlation matrix at time t can also be derived in terms of the initial correlations. Its eigendecomposition can be viewed as a particular variant of proper orthogonal decomposition and provides the most statistically prevalent structures, which serve as the statistical analogue of the left singular vectors of the evolution matrix.

Quantifying the likelihood that a disturbance grows beyond a particular level requires knowledge of the probability density function (p.d.f.) of the energy amplification. Whereas the mean energy amplification depends only on the correlation matrix of the incoming disturbances, the entire distribution of incoming disturbances is needed to calculate the p.d.f. of the growth. Moreover, there is no general formula relating the two. However, we observe empirically that the p.d.f. is nearly exponential, and this leads to an approximation strategy for it. We use the approximate p.d.f. to derive accurate confidence bounds on the growth, i.e. energy levels which $p\%$ of the disturbances do not exceed, for some desired p . The exponential behaviour of the p.d.f. also means that if G^{mean} is significantly below G^{opt} , it is extremely unlikely for an initial disturbance to achieve near- G^{opt} growth. In particular, if G^{opt} is k times larger than G^{mean} , the probability of near-optimal growth is e^{-k} due to the exponential p.d.f.

Throughout the paper, we demonstrate the statistical framework using plane Poiseuille flow. Equipped with a statistical lens, numerous observations readily emerge. At each wavenumber pair (α, β) , the correlation length in the wall-normal direction has a dramatic impact on G^{mean} , with correlation lengths of the order of the channel half-height growing to nearly half of G^{opt} . If, however, the correlation length is short compared with the channel half-height, G^{mean} can be orders of magnitude smaller than G^{opt} ; G^{opt} and G^{mean} achieve their maximum values at similar locations in wavenumber space, but the peak is substantially narrower in α for G^{mean} . This indicates that three-dimensional disturbances, ones which contain a range of wavenumbers, further undershoot G^{opt} . In the three-dimensional case, G^{mean} is a function of the three-dimensional correlation matrix. We observe that when this correlation is isotropic, G^{mean} is roughly 2.5% of G^{opt} at $Re = 1000$. Surprisingly, we find that G^{mean} scales nearly linearly with Re , so the gap

between it and G^{opt} widens with increasing Reynolds number. Therefore, G^{opt} increasingly overstates the growth of random disturbances.

Even considering disturbances near the optimal wavenumber pair ($\alpha = 0, \beta = 2$), the probability of exceeding certain levels of growth can be extremely low. We show that the distribution of energy is nearly exponential, i.e. the probability of exceeding a particular energy level decays exponentially. Therefore, if G^{mean} is relatively small relative to G^{opt} , there is little chance of observing growth of the order of the optimal value. For a correlation length of one fourth the channel half-height, fewer than 0.01 % of disturbances achieve $G^{opt}/2$ growth for $Re = 1000$.

The combined effects of the non-normality of the linearized Navier–Stokes operator and randomness have been analysed before. In particular, Farrell & Ioannou (1993, 1994, 1996) and later Fontane, Brancher & Fabre (2008) considered the linearized Navier–Stokes equations forced continuously by white-in-time noise with some spatial correlation. They showed that the expected energy, once statistical stationarity is reached, can be obtained by solving a Lyapunov equation involving the linearized Navier–Stokes operator. Our study is distinct from this work in two ways. First, rather than using a continuously forced model, we use the physical model of transient growth, wherein the linearized equations are impulsively disturbed at $t = 0$, and the disturbance evolves without further forcing. Second, we explore the effect of different initial disturbance statistics, whereas Farrell & Ioannou (1993, 1994, 1996) assume a temporally white forcing and do not assess the impact of different forcing spatial correlations on the expected energy of the disturbance. In [Appendix A](#), we detail a statistical formulation for the continuously forced case that is analogous to the framework we present for transient growth. There, one has to supply the spatio-temporal correlation of the forcing instead of the spatial correlation of the initial disturbance, and the white-noise model in Farrell & Ioannou (1993, 1994, 1996) and Fontane *et al.* (2008) is recovered as a special case. One quantitative comparison can be made between the white-noise model and our statistical formulation of transient growth: the white-noise model gives a Reynolds number scaling between $Re^{1.5}$ and Re^3 , depending on the wavenumber, while our results show a scaling of Re^1 for disturbances that are broadband in wavenumber and Re^2 for single wavenumber pairs. The present work is also different from what has been called statistical stability (Malkus 1956; Markeviciute 2022). That work is concerned with the stability of the statistical state of turbulent flow, whereas our study investigates the statistics of transient growth.

The remainder of the paper is organized as follows. In § 2, plane Poiseuille flow and the numerics used to perform the calculations are described. Section 3 gives a review of transient growth. In § 4, we derive a formula for the mean energy amplification and compare it with the optimal growth for Poiseuille flow, first for disturbances at one pair of wavenumbers, then for disturbances containing a range of wavenumbers. We investigate the p.d.f. of the growth and detail an accurate approximation strategy for it in § 5. Finally, in § 6, we conclude the paper.

2. Flow description and numerics

Plane Poiseuille flow is the steady, laminar flow between two plates separated by $2h$ in the y direction. The flow is in the x direction, and the plates are infinite in both the streamwise (x) and spanwise (z) directions. It is driven by a constant pressure gradient and the streamwise velocity field is given by

$$U(y) = -\frac{1}{2\mu} \frac{dp}{dx} (h^2 - y^2), \quad (2.1)$$




Figure 1. Schematic of Poiseuille flow. The red waves represent disturbances with particular wavenumbers.

where p is the pressure and μ is the dynamic viscosity. The flow is then non-dimensionalized by the channel half-height and the centreline velocity. Because the governing equations and base flow are homogenous in x and z , it is convenient to take the Fourier transform of disturbances to the base flow in these directions. The associated wavenumbers in the streamwise and spanwise directions are denoted α and β , respectively. For example, the transformed wall-normal velocity is

$$\hat{v}(y, \alpha, \beta) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} v(x, y, z) \exp(-i(\alpha x + \beta z)) \, dx \, dz. \quad (2.2)$$

The physical set-up is shown in [figure 1](#). Employing the usual velocity–vorticity formulation of the linearized Navier–Stokes equations yields the following equations for the evolution of disturbances ([Reddy & Henningson 1993](#)):

$$\frac{\partial}{\partial t} \begin{bmatrix} \hat{v} \\ \hat{\eta} \end{bmatrix} = -i \begin{bmatrix} \mathcal{L}_{OS} & 0 \\ \mathcal{L}_C & \mathcal{L}_{SQ} \end{bmatrix} \begin{bmatrix} \hat{v} \\ \hat{\eta} \end{bmatrix}. \quad (2.3)$$

The Orr–Sommerfeld, cross-term and Squire operators are

$$\mathcal{L}_{OS} = - \left(\frac{\partial^2}{\partial y^2} - k^2 \right)^{-1} \left[\frac{1}{iRe} \left(\frac{\partial^2}{\partial y^2} - k^2 \right)^2 - \alpha U \left(\frac{\partial^2}{\partial y^2} - k^2 \right) + \alpha U'' \right], \quad (2.4a)$$

$$\mathcal{L}_C = \beta U', \quad (2.4b)$$

$$\mathcal{L}_{SQ} = \alpha U - \frac{1}{iRe} \left(\frac{\partial^2}{\partial y^2} - k^2 \right). \quad (2.4c)$$

Above, all quantities are non-dimensionalized, and $U = U(y)$ is the base flow, $\hat{\eta}$ is the transformed wall-normal vorticity, $k^2 = \alpha^2 + \beta^2$ is the squared wavevector magnitude and $(\cdot)'$ indicates a wall-normal derivative $\partial/\partial y$. We use the code provided in [Schmid & Henningson \(2001\)](#), which uses a Chebyshev discretization of the linearized Navier–Stokes equations (2.3) ([Herbert 1977](#); [Reddy & Henningson 1993](#)). All norms presented in our numerical results are based on the kinetic energy of a disturbance. It can be shown, by using incompressibility and Parseval’s theorem, that the energy of a disturbance in the transformed velocity–vorticity coordinates is ([Gustavsson 1986](#))

$$e = \frac{1}{4\pi^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{1}{2k^2} \int_{-1}^1 \left\| \frac{\partial}{\partial y} \hat{v} \right\|^2 + k^2 \|\hat{v}\|^2 + \|\hat{\eta}\|^2 \, dy \, d\alpha \, d\beta. \quad (2.5)$$

3. Optimal transient growth

Here, we review the linear effects responsible for transient growth in a system with all negative eigenvalues. For a more thorough review, see Schmid (2007). Expressing the Navier–Stokes equations as

$$\dot{\tilde{q}}(\mathbf{x}, t) = \mathcal{N}(\tilde{q}(\mathbf{x}, t)), \quad (3.1)$$

a steady solution $\bar{q}(\mathbf{x})$ satisfies $\mathcal{N}(\bar{q}(\mathbf{x})) = \mathbf{0}$. Although their size is likely small, disturbances to the base flow are inevitable. Denoting these disturbances as $\mathbf{q}(\mathbf{x}, t) = \tilde{q}(\mathbf{x}, t) - \bar{q}(\mathbf{x})$, their dynamics are analysed by linearizing around the base flow

$$\dot{\mathbf{q}}(\mathbf{x}, t) = \mathcal{N}(\bar{q}(\mathbf{x}) + \mathbf{q}(\mathbf{x}, t)) \approx \mathbf{A}\mathbf{q}(\mathbf{x}, t), \quad (3.2)$$

where \mathbf{A} is the Jacobian around the base flow

$$\mathbf{A} = \left. \frac{\partial \mathcal{N}}{\partial \mathbf{q}} \right|_{\bar{q}}. \quad (3.3)$$

The problem is discretized as

$$\dot{\mathbf{q}}(t) = \mathbf{A}\mathbf{q}(t), \quad (3.4)$$

where $\mathbf{q}(t) \in \mathbb{R}^N$ is the discretized state vector describing the disturbance.

The solution to (3.5) is

$$\mathbf{q}(t) = \mathbf{M}_t \mathbf{q}(0), \quad (3.5)$$

where the evolution operator is the matrix exponential

$$\mathbf{M}_t = \exp[\mathbf{A}t]. \quad (3.6)$$

If all of the eigenvalues of the linear operator \mathbf{A} have a negative real part, then the linear system is stable in the sense that the norm of any infinitesimal disturbance will eventually decay, i.e. $\lim_{t \rightarrow \infty} \|\mathbf{q}(t)\| = 0$. This sense of stability, usually referred to as modal stability, is mathematically powerful – it is a property of the system, not of any particular disturbance. Assuming the linear approximation is valid, if the eigenvalues are negative, any disturbance decays eventually, but if there is a positive eigenvalue, any disturbance arising in a physical scenario will have a non-zero projection onto the associated eigenvector, and will thus grow exponentially.

The theory of transient growth offers the additional insight that, even if all the eigenvalues are stable, if \mathbf{A} is non-normal, i.e. its eigenvectors are non-orthogonal, the decay need not be monotonic. The eigenvectors summed together to construct an initial disturbance may mostly cancel each other initially, but because they vanish at different rates, after some time, there may no longer be cancellation, which leads to growth of the disturbance. Physically, this can be viewed as a constructive interference of certain combinations of the eigenvectors. The linear operators arising in fluid systems, especially in shear flows, can be highly non-normal (Trefethen *et al.* 1993). The ability for these systems to produce growth is quantified in the literature by the maximal amplification that a disturbance may undergo

$$G^{opt}(t) \equiv \max_{\|\mathbf{q}(0)\|=1} \|\mathbf{q}(t)\|^2. \quad (3.7)$$

This quantity is usually referred to simply as G . Here, we have termed it G^{opt} to specify that it is the optimal growth among all possible initial disturbances and to distinguish it from G^{mean} , which will arise later in the paper. Its peak in time is referred to in this paper

as G_{max}^{opt} (usually referred to simply as G_{max}). The norm $\|\cdot\|$ is based on the kinetic energy of the disturbance and can be written

$$e(\mathbf{q}) = \|\mathbf{q}\|^2 = \mathbf{q}^* \mathbf{W} \mathbf{q}. \quad (3.8)$$

Here, $(\cdot)^*$ denotes Hermitian conjugation, \mathbf{W} is a weight matrix (required to be Hermitian and positive-definite), and we make frequent use of the decomposition $\mathbf{L}^* \mathbf{L} = \mathbf{W}$. For later use, the inner product that induces the norm is $\langle \mathbf{q}_1, \mathbf{q}_2 \rangle = \mathbf{q}_2^* \mathbf{W} \mathbf{q}_1$. It can be shown that the optimal growth may be written (Reddy & Henningson 1993)

$$G^{opt}(t) = \sigma_1^2(\mathbf{L} \mathbf{M}_t \mathbf{L}^{-1}), \quad (3.9)$$

where $\sigma_1^2(\cdot)$ returns the first (squared) singular value of the argument. The structures that undergo the most growth up to time t and the structures resulting from the amplification may also be obtained via the singular value decomposition (SVD) of the weighted evolution operator

$$\mathbf{L} \mathbf{M}_t \mathbf{L}^{-1} = \tilde{\mathbf{U}} \tilde{\Sigma} \tilde{\mathbf{V}}^*. \quad (3.10)$$

The optimal output and input modes are recovered as $\mathbf{U} = \mathbf{L}^{-1} \tilde{\mathbf{U}}$ and $\mathbf{V} = \mathbf{L}^{-1} \tilde{\mathbf{V}}$, respectively. The first column of \mathbf{V} is the initial disturbance that grows by $G^{opt}(t)$, and the first column of \mathbf{U} is the structure that results.

The largest initial growth rate experienced by any disturbance can be expressed in terms of the optimal growth as

$$a^{opt} = \left. \frac{d}{dt} G^{opt}(t) \right|_{t=0}. \quad (3.11)$$

By expanding the matrix exponential to first-order terms in t , it is easily shown that this optimal growth rate is given by the numerical abscissa (Trefethen & Embree 2005)

$$a^{opt} = \kappa_1(\mathbf{L} \mathbf{A} \mathbf{L}^{-1} + (\mathbf{L} \mathbf{A} \mathbf{L}^{-1})^*), \quad (3.12)$$

where $\kappa_1(\cdot)$ returns the first eigenvalue of the argument.

So long as the disturbance remains small enough, the linear approximation (3.2) remains valid, and the disturbance will decay to zero. However, if the growth is large enough, it can elevate a disturbance from the regime where linearity governs to one where nonlinear effects are relevant. These nonlinear effects can in turn lead the flow away from the base state, eventually causing transition. The growth can indeed be quite large, owing to the severe non-normality in the linearized Navier–Stokes operator in shear flows. Figure 2(a) shows $G^{opt}(t)$ for various streamwise and spanwise wavenumbers in plane Poiseuille flow at $Re = 1000$. For $\alpha = 0$, $\beta = 2$, G_{max}^{opt} is nearly 200. Figure 2(b) shows G_{max}^{opt} for a range of wavenumbers. Streamwise-elongated structures (small α) are capable of larger growth than shorter structures (larger α). The peak in wavenumber space is at $\alpha = 0$, $\beta = 2.04$, so structures of finite spanwise (z) length experience the most growth.

To motivate the remainder of this paper, we show 1000 random trajectories along with $G^{opt}(t)$ at $Re = 1000$, $\alpha = 0$, $\beta = 2$ in figure 3. Indeed, G^{opt} bounds the trajectories, but, notably, they all substantially undershoot it. The details of the distribution used to




Figure 2. Optimal gains for plane Poiseuille flow at $Re = 1000$. (a) The maximal gain over all initial disturbances $G^{opt}(t)$ for various choices of wavenumbers. (b) Maximal gain, also maximized over time for a range of streamwise and spanwise wavenumbers α, β .




Figure 3. Value of $G^{opt}(t)$ for $\alpha = 0, \beta = 2$ along with 1000 random trajectories.

generate figure 3 are given in § 5. In what follows, we derive formulae to describe the statistics of the growth and demonstrate them on plane Poiseuille flow, recording our observations.

4. Expected energy amplification

In light of figure 3, an obvious question is: How much energy, on average, do the amplified disturbances achieve? We derive a formula for the mean energy of the amplified disturbances in terms of the correlation matrix of the initial disturbances. The expected energy divided by the expected initial energy is termed G^{mean} . We elaborate on the difference between this and the expected value of the ratio of these energies at the end of the following subsection, but, in short, G^{mean} is more physically meaningful, produces a simpler mathematical result, and requires less *a priori* knowledge of the initial disturbances.

Just as in the standard treatment of transient growth, the physical model that we consider consists of the discretized base flow $\bar{\mathbf{q}}$, which is impulsively perturbed at $t = 0$ by $\mathbf{q}(0)$. As before, the disturbance $\mathbf{q}(0)$ may represent the entire (three-dimensional) flow in space, the flow at a particular pair of wavenumbers or the flow at a particular location in the streamwise direction. The evolution of the disturbances is governed by the Navier–Stokes equations linearized around the base flow, so (3.5) holds. However, our statistical framework differs from the standard treatment of transient growth in that the disturbance $\mathbf{q}(0)$ is now a random variable with some distribution, and we study the statistics of the disturbance after some time $\mathbf{q}(t)$. In particular, we are interested in the energy of the growing disturbance in comparison with that of the initial disturbance.

Experimenting with various choices of the initial correlation for plane Poiseuille flow reveals that the expected energy can be substantially smaller than G_{max}^{opt} . This is especially true when the correlation length is short relative to the channel half-height. Furthermore, G_{max}^{mean} drops off more rapidly with larger α than does G_{max}^{opt} , which causes the mean energy amplification for three-dimensional disturbances to be quite small unless their energy is focused sharply at $\alpha = 0$. Surprisingly, we observe that, for isotropically correlated initial disturbances, the mean energy amplification scales near-linearly with Re , in contrast to the quadratic scaling of G_{max}^{opt} .

4.1. Theory

4.1.1. The quantity G^{mean}

For simplicity, we omit the weight matrix in the derivations (by setting it to the identity), reporting the formulae with it at the end, so $e(\mathbf{q}(t)) = \mathbf{q}^*(t)\mathbf{q}(t)$. The energy may alternatively be written as the trace of the outer product

$$e(\mathbf{q}(t)) = \text{Tr}\{\mathbf{q}(t)\mathbf{q}^*(t)\}, \tag{4.1}$$

because the diagonals of $\mathbf{q}(t)\mathbf{q}^*(t)$ are the terms summed in the inner product. In terms of the evolution operator, (4.1) becomes

$$e(\mathbf{q}(t)) = \text{Tr}\{\mathbf{M}_t\mathbf{q}(0)\mathbf{q}^*(0)\mathbf{M}_t^*\}. \tag{4.2}$$

The expected value of this expression gives the expected energy of the amplified disturbances

$$\mathbb{E}[e(\mathbf{q}(t))] = \mathbb{E}[\text{Tr}\{\mathbf{M}_t\mathbf{q}(0)\mathbf{q}^*(0)\mathbf{M}_t^*\}]. \tag{4.3}$$

The expectation commutes with the trace and evolution matrices, giving

$$\mathbb{E}[e(\mathbf{q}(t))] = \text{Tr}\{\mathbf{M}_t\mathbb{E}[\mathbf{q}(0)\mathbf{q}^*(0)]\mathbf{M}_t^*\}. \tag{4.4}$$

The expectation of the outer product of the initial disturbances is their correlation matrix

$$\mathbf{C}_{00} = \mathbb{E}[\mathbf{q}(0)\mathbf{q}^*(0)], \tag{4.5}$$

so the expected energy of the growing disturbance is expressed in terms of the correlations of the initial disturbances

$$\mathbb{E}[e(\mathbf{q}(t))] = \text{Tr}\{\mathbf{M}_t\mathbf{C}_{00}\mathbf{M}_t^*\}. \tag{4.6}$$

A metric for the expected growth of the disturbances, which we term G^{mean} , is provided by the ratio of the expected energy and initial energy

$$G^{mean}(t) \equiv \frac{\mathbb{E}[e(\mathbf{q}(t))]}{\mathbb{E}[e(\mathbf{q}(0))]} = \frac{\text{Tr}\{\mathbf{M}_t\mathbf{C}_{00}\mathbf{M}_t^*\}}{\text{Tr}\{\mathbf{C}_{00}\}}. \tag{4.7}$$

This quantity is not the same as the expected value of the growth; this difference is discussed at the end of this subsection. If a weight matrix $\mathbf{W} = \mathbf{L}^*\mathbf{L}$ is used to define

the energy, then (4.7) becomes

$$G^{mean}(t) = \frac{\text{Tr}\{\mathbf{L}\mathbf{M}_t\mathbf{C}_{00}\mathbf{M}_t^*\mathbf{L}^*\}}{\text{Tr}\{\mathbf{L}\mathbf{C}_{00}\mathbf{L}^*\}}. \tag{4.8}$$

The value of $G^{opt}(t)$ is given in terms of the SVD of the evolution operator. To express $G^{mean}(t)$ in a similar manner, we make use of the fact that the trace of a matrix is the sum of its eigenvalues and that the eigenvalues of $\mathbf{B}\mathbf{B}^*$ are the squared-singular values of \mathbf{B} for any matrix \mathbf{B} . Using these two facts, (4.7) can be written

$$G^{mean}(t) = \frac{\sum_{i=1}^N \sigma_i^2(\mathbf{M}_t\mathbf{B})}{\sum_{i=1}^N \sigma_i^2(\mathbf{B})}, \tag{4.9}$$

where \mathbf{B} is defined by the factorization $\mathbf{C}_{00} = \mathbf{B}\mathbf{B}^*$. In the case of a weight matrix, (4.9) becomes

$$G^{mean}(t) = \frac{\sum_{i=1}^N \sigma_i^2(\mathbf{L}\mathbf{M}_t\mathbf{B})}{\sum_{i=1}^N \sigma_i^2(\mathbf{L}\mathbf{B})}. \tag{4.10}$$

Upper and lower bounds for $G^{mean}(t)$ for any possible initial correlation can be obtained by setting \mathbf{C}_{00} to the outer product of the first input mode with itself and last input mode with itself, i.e. $\mathbf{v}^1\mathbf{v}^{1*}$ and $\mathbf{v}^N\mathbf{v}^{N*}$, respectively, yielding the bounds $\sigma_1^2(\mathbf{L}\mathbf{M}_t\mathbf{L}^{-1})$ and $\sigma_N^2(\mathbf{L}\mathbf{M}_t\mathbf{L}^{-1})$. Notably, the upper bound is $G^{opt}(t)$. In the case that the disturbances are white in space, i.e. $\mathbf{C}_{00} = \mathbf{W}^{-1}$, the resulting $G^{mean}(t)$ is the mean-squared singular value of the weighted evolution operator $\mathbf{L}\mathbf{M}_t\mathbf{L}^{-1}$.

The quantity G^{mean} , defined in (4.7), is the ratio of the expected energy of the disturbance at time t to its expected initial energy. This is distinct from the expected ratio of energy, $\mathbb{E}[e(\mathbf{q}(t))/e(\mathbf{q}(0))]$. Physically, the ratio of expected energies is the more salient quantity because whether a particular disturbance leads to transition depends on its final energy (and shape), not on the growth it underwent. In figure 3, this ratio of expected energies is the mean of the grey curves (at each time). The expected ratio of energies would come from dividing each disturbance by its initial energy, then taking the average, but this inappropriately weights the growth of smaller initial disturbances equal to that of larger ones. Mathematically, the ratio of expected energies is the easier quantity to work with because it depends only on the correlation matrix of the initial disturbances, as shown in (4.7), while the expected ratio of energies depends on the entire distribution of the initial disturbances. If there is no variation in the size of initial disturbances, i.e. if they live on an N -dimensional sphere, the two quantities are the same. More generally, the quantities are the same in the case that the distribution of initial disturbances is separable in radius and direction, as is proven in Appendix B.

Analogous to the numerical abscissa a^{opt} , we define the mean initial growth rate

$$a^{mean} \equiv \left. \frac{d}{dt} G^{mean}(t) \right|_{t=0}. \quad (4.11)$$

This derivative can be calculated by expanding the evolution operator to first order

$$a^{mean} = \left. \frac{d}{dt} \frac{\text{Tr}\{\mathbf{L}(\mathbf{I} + t\mathbf{A})\mathbf{C}_{00}(\mathbf{I} + t\mathbf{A}^*)\mathbf{L}^*\}}{\text{Tr}\{\mathbf{L}\mathbf{C}_{00}\mathbf{L}^*\}} \right|_{t=0}, \quad (4.12)$$

where \mathbf{I} is the identity. Dropping the quadratic term and evaluating the derivative gives

$$a^{mean} = \frac{\text{Tr}\{\mathbf{L}\mathbf{A}\mathbf{C}_{00}\mathbf{L}^* + \mathbf{L}\mathbf{C}_{00}\mathbf{A}^*\mathbf{L}^*\}}{\text{Tr}\{\mathbf{L}\mathbf{C}_{00}\mathbf{L}^*\}}. \quad (4.13)$$

Finally, leveraging the Hermiticity of the correlation matrix

$$a^{mean} = 2 \frac{\text{Tr}\{\text{Re}(\mathbf{L}\mathbf{A}\mathbf{C}_{00}\mathbf{L}^*)\}}{\text{Tr}\{\mathbf{L}\mathbf{C}_{00}\mathbf{L}^*\}}, \quad (4.14)$$

where $\text{Re}(\cdot)$ returns the real part of the argument. The upper bound for this quantity is a^{opt} , which is positive if (and only if) $G^{opt} > 1$, but we have never observed a^{mean} to be positive in our numerical experiments. Indeed, we have never observed a randomly chosen disturbance initially grow.

4.1.2. Correlation and dominant structures

The statistics of the initial disturbances can also be used to augment prediction of the structures that arise from the linear amplification by the evolution operator. Removing the trace from (4.6) gives a formula for the correlation matrix of the disturbance at time t

$$\mathbf{C}_{tt} \equiv \mathbb{E}[q(t)q(t)^*] = \mathbf{M}_t \mathbf{C}_{00} \mathbf{M}_t^*. \quad (4.15)$$

The dominant flow structures at time t are the eigenvectors of this correlation matrix (multiplied by a weight if desired)

$$\mathbf{C}_{tt} \mathbf{W} \Phi_t = \Phi_t \Lambda_t. \quad (4.16)$$

The columns ϕ_t^k of Φ_t are orthogonal in the weighted inner product, i.e. $\langle \phi_t^i, \phi_t^j \rangle = \delta_{ij}$. This can be thought of as a particular variant of proper orthogonal decomposition (POD) (Lumley 1967, 1970; Sirovich 1987) in which the data consist of an ensemble of realizations of the disturbances at a specific time t rather than a single time series. The eigenvalues are non-negative, owing to the semi-positive definiteness of the correlation matrix, and represent the expected energy of each structure. More precisely, the k th eigenvalue

$$\lambda_t^k = \mathbb{E}[|\langle q(t), \phi_t^k \rangle|^2], \quad (4.17)$$

is the expected energy of the projection of the disturbance onto the k th mode ϕ_t^k . The eigenvalues sum to the total expected energy, so

$$G^{mean}(t) = \frac{\sum_i \lambda_t^i}{\sum_i \lambda_0^i}. \quad (4.18)$$

Therefore, the eigenvalues quantify the expected contribution of each mode to the growth of the disturbance.

The average energy of the disturbance captured by any structure can be quantified (Frame & Towne 2023) by

$$\epsilon(\boldsymbol{\psi}) = \mathbb{E}[|\langle \boldsymbol{q}(t), \boldsymbol{\psi} \rangle|^2]. \quad (4.19)$$

The first POD mode maximizes this quantity (over normalized modes), and the latter modes maximize it with the constraint that they are orthogonal to all previous ones. For a more thorough review of POD, see Rowley & Dawson (2017), Taira *et al.* (2017) or Towne, Schmidt & Colonius (2018).

The POD modes offer an alternative to the output modes of the evolution matrix for describing the structures that emerge from the linear amplification. The POD modes are the most energetic structures, while the output modes are the modes resulting from the greatest amplification by the evolution operator. In the case that the initial correlation \mathbf{C}_{00} is white with respect to the weight, the POD modes are equivalent to the output modes, i.e.

$$\mathbf{C}_{00} = \mathbf{W}^{-1} \implies \boldsymbol{\Phi} = \mathbf{U}. \quad (4.20)$$

This result is analogous to the relationship between resolvent modes and spectral POD modes established by Towne *et al.* (2018). Of course, the initial correlation is unlikely to be white in a real flow, so it is advantageous to use knowledge of the incoming statistics to augment the prediction of these structures.

In the remainder of this section, we experiment with different choices of \mathbf{C}_{00} for Poiseuille flow and record our observations. We are not aware of any previous studies on the initial correlations of disturbances within Poiseuille flow. Indeed, in the context of temporal stability, \mathbf{C}_{00} is not the usual correlation one would measure in an experiment. In an experiment, one is measuring not only the initial disturbances but also disturbances that have already been evolved by the linear dynamics. This makes it difficult to distinguish between the initial disturbances and the time-evolved ones, which means that one cannot unambiguously determine the correlation matrix of the initial disturbances. Additionally, the nature of the disturbances and quantities, such as their initial correlations, are certainly sensitive to the specifics of the flow set-up. For example, the disturbances generated by vibrations of the boundary are likely substantially different from those caused by surface roughness. Providing a model for the correlations of the initial disturbance is not the topic of this paper, and we do not claim that the choices made below are necessarily reflective of the physics in Poiseuille flow. However, the trends that emerge, e.g. that longer correlation lengths lead to more growth and that G^{mean} is substantially smaller than G^{opt} and grows linearly when three-dimensional effects are accounted for, are not specific to our choice of the correlation, and therefore give physical insight despite the current lack of an accurate model for the correlations.

4.2. Numerical experiments with disturbances at a single wavenumber pair

Most studies of transient growth in flows with homogeneous directions take the Fourier transform in these directions and calculate the transient growth for disturbances consisting of a single pair of streamwise and spanwise wavenumbers. Here, we perform the analogous analysis for G^{mean} in Poiseuille flow. The correlation at a particular α and β can be written

$$\hat{\mathbf{C}}_{00}(y_1, y_2; \alpha, \beta) = \begin{bmatrix} \hat{\mathbf{C}}_{00}^{vv}(y_1, y_2; \alpha, \beta) & \hat{\mathbf{C}}_{00}^{v\eta}(y_1, y_2; \alpha, \beta) \\ \hat{\mathbf{C}}_{00}^{\eta v}(y_1, y_2; \alpha, \beta) & \hat{\mathbf{C}}_{00}^{\eta\eta}(y_1, y_2; \alpha, \beta) \end{bmatrix}, \quad (4.21)$$

where the diagonal terms are the autocorrelations of wall-normal velocity and wall-normal vorticity, and the off-diagonal terms are the cross-correlations between these two variables.




Figure 4. Value of $G^{mean}(t)$ (solid) and $G^{opt}(t)$ (dashed) for various wavenumbers at $Re = 1000$. The mean energy amplification is substantially higher for the longer correlation length $\lambda^{-1} = 1$ than for the shorter one $\lambda^{-1} = 5$.

It can be shown analytically that, for a disturbance to experience large growth, its initial energy should be concentrated in its wall-normal velocity rather than wall-normal vorticity (Gustavsson 1991), and we have observed this property to persist within the statistical framework. Therefore, we choose only the vertical velocity autocorrelation to be non-zero and take it to be Gaussian in the wall-normal direction with correlation length λ , i.e.

$$\hat{\mathbf{C}}_{00}^{vv}(y_1, y_2; \alpha, \beta) = \mathbb{E}[\hat{v}(y_1, \alpha, \beta)\hat{v}^*(y_2, \alpha, \beta)] = A \exp\left[-\frac{(y_1 - y_2)^2}{\lambda^2}\right]. \quad (4.22)$$

The normalization A has no impact on G^{mean} because this constant affects the expected energy of the amplified disturbances and that of the initial ones equally. In our numerics, it is chosen so that when the initial correlation is discretized in y , its trace is unity.

4.2.1. Mean energy amplification for a single wavenumber pair

Figure 4 shows $G^{mean}(t)$ (solid) for various wavenumbers and $G^{opt}(t)$ (dashed) for the same wavenumbers, both as functions of time for $Re = 1000$. Whether the mean is of the same order as the maximum depends on the characteristics of the correlations of the initial disturbances. We refer to the peak of $G^{mean}(t)$ in time as G_{max}^{mean} . For the relatively long correlations in (a), G_{max}^{mean} is roughly half G_{max}^{opt} for the most amplified wavenumbers, while for the shorter correlation length (b), the ratio is closer to one tenth. Figure 5 shows the first time unit of $G^{mean}(t)$ using the same parameters as figure 4. Despite the fact that G^{mean} grows to be relatively large, it initially decays sharply for all wavenumbers. The initial decay rate can be calculated with (4.14).

Figure 6(a) shows G_{max}^{mean} for a range of λ^{-1} at $Re = 1000$. The correlation length λ greatly impacts the mean energy amplification, with longer correlation lengths corresponding to more growth and shorter ones to less growth. It is likely this trend is explained by the fact that short-wavelength (in y) disturbances are quickly dissipated by viscosity before they can extract energy from the mean shear (McKeon 2017).




Figure 5. Value of $G^{mean}(t)$ for short times. A steep decay is observed initially even in cases where G^{mean}_{max} is relatively high. The initial growth (or decay) rate is a^{mean} , given in (4.14).




Figure 6. The effect of correlation length on G^{mean}_{max} for Poiseuille flow at $Re = 1000$. Panel (a) shows $G^{mean}(t)$ maximized over time vs inverse correlation length for various streamwise and spanwise wavenumbers. More coherent disturbances (large λ) tend to grow more, but there is a non-infinite optimum. (b) The time at which G^{mean} is maximized vs inverse correlation length. The maximum time does not vary much with λ but does with α and β , with shorter wavelengths corresponding to an earlier maximization time. The maximization time drops to zero when the correlation becomes such that $G^{mean}(t)$ never exceeds one.

A related point is worth discussing. One may be interested in G^{mean} resulting from a spatially white correlation matrix, i.e. a correlation matrix that is a multiple of the identity. A white correlation may seem appealing at first glance because it represents a ‘neutral’ choice for the initial statistics. However, this choice yields grid-dependent results for the growth. In fact, with twice as many grid points (and half the grid spacing), we observed that G^{mean} decreases by a factor of roughly two. Upon some reflection, this grid dependence is not a surprise for two reasons. First, the physical meaning of the discretized representation of white noise on the coarse and fine grids is different. Sample disturbances on the grid are jagged, with the value at each grid point uncorrelated from the values at adjacent grid points. When the grid spacing is halved, samples become twice as jagged,




Figure 7. Dependence of G^{mean} on the streamwise and spanwise wavenumbers at $Re = 1000$ for different correlation lengths. The shape is similar to the contour of G_{max}^{opt} at the same Reynolds number (figure 2), but, notably, the support in α is substantially narrower for G^{mean} . This indicates that the energy of the disturbance must be quite concentrated at the large-growth wavenumbers to achieve significant growth.

hence, the physical meaning of the discretized representation of white noise depends on the grid used to represent it. Second, the white noise on the finer grid concentrates a greater fraction of the energy at shorter scales because white noise is uniform in its wavenumber content among wavenumbers representable on the grid. Figure 6(a) shows that very short scales do not grow, so one should expect that G^{mean} for white noise on the finer grid will be smaller because less of the initial energy is concentrated at scales that grow. For meaningful, i.e. grid-independent, results, one must use a correlation matrix that can be fully resolved, and we stress that we have done this throughout the paper.

Figure 6(b) shows the time at which G^{mean} is maximized. This time is relatively independent of the correlation length, but changes substantially with the wavenumber pair. The substantial wavenumber dependence of the maximization time is important when considering the growth of disturbances that are broadband in wavenumber. The discontinuity is not an artefact. For λ^{-1} just before the discontinuity, $G^{mean}(t)$ increases very gradually near $t = 0$, reaching its peak at $t \approx 80$. For λ^{-1} just after the discontinuity, $G^{mean}(t)$ decreases very gradually near $t = 0$, so the maximum growth occurs at $t = 0$ and the maximization time suddenly drops from near 80 to 0.

The wavenumber dependence is further explored in figure 7. Figure 7(a) shows this dependence for $\lambda^{-1} = 1$, which is near the peak for the maximally amplified wavenumber in figure 6(a). The location of the peak in wavenumber space is near that of G_{max}^{opt} seen in figure 2; however, G_{max}^{mean} decays much more rapidly with α than does G_{max}^{opt} . This indicates that to achieve large-scale growth, the energy of a disturbance must be narrowly concentrated in wavenumber space at the values that experience large growth. As shown in §4.4, this severely limits the mean energy amplification of fully three-dimensional disturbances.

Figure 8 shows G_{max}^{mean} at $\alpha = 0$, $\beta = 2$ for a range of Reynolds numbers and correlation lengths. Similar dependence on correlation length is observed at all Reynolds numbers with the peak occurring when the correlation length is roughly the channel half-height. The value of G_{max}^{opt} is known to scale quadratically with Reynolds number (for small values of αRe) (Gustavsson 1991). In figure 9, we show the scaling of G_{max}^{mean} with Reynolds




Figure 8. Value of G_{max}^{mean} at $\alpha = 0, \beta = 2$ as a function of Reynolds number and inverse correlation length λ^{-1} . Similar behaviour is observed over the range of Re .




Figure 9. Scaling of G_{max}^{mean} for disturbances at a single wavenumber pair for a variety of correlation lengths at $\alpha = 0, \beta = 2$. The scaling appears quadratic (the grey dashed line is Re^2), matching that of G_{max}^{opt} .

number for a variety of correlation lengths at $\alpha = 0, \beta = 2$. These appear to obey the same scaling.

4.2.2. Dominant structures for a single wavenumber pair

Now we examine the structures that emerge in Poiseuille flow at a single wavenumber pair, as described in § 4.1.2. The key question is: To what extent do the output modes resemble the principal components of the correlation matrix, i.e. the POD modes? The former are the structures resulting from the largest amplification by the evolution operator (see (3.10)), whereas the latter use the statistics of the initial disturbances to inform which structures are most energetic (see (4.16)).

Figure 10 shows the evolution of the correlations and the vorticity component of their POD modes for $\alpha = 0, \beta = 2$. We impose the initial correlation to be of the form in (4.22) with $\lambda = 1$, so all the energy is initially in the velocity. The evolution operator rapidly shifts this energy to the vorticity, and two counter-rotating vortices emerge. The first POD mode (blue) reflects this with two peaks of opposite sign at the peaks in the vorticity correlation. The velocity component of the POD mode is not plotted because it rapidly decays to 0. Notably, the first output mode quickly becomes nearly identical to the first POD mode, despite the former not depending on the initial correlation. Indeed, for this




Figure 10. Evolution of correlations and their POD modes for $Re = 1000$, $\alpha = 0$, $\beta = 2$ and $\lambda = 1$. The initial energy is imposed to be in the velocity (a–d), but it is quickly shifted to vorticity (e–h). The first POD mode (blue dashed) and first output mode (red) quickly become similar (i–l). For these parameters, $G^{mean}(t)$ peaks near $t = 80$ (see figure 4).

wavenumber pair, the first few POD modes from different initial correlation matrices quickly become similar to one another and to the first few left singular values of the evolution matrix. There is only moderate gain separation in the singular values, so the similarity between the modes is surprising.

Figure 11 compares the modes more thoroughly. For each wavenumber pair, figure 11(a) shows the average energy captured (see (4.19)) by the first output mode at the peak time of G^{mean} as a fraction of that captured by the first POD mode at the same time. The energy captured by the POD mode is the maximum possible, so a value near unity indicates that the output mode is very effective, while a value near zero indicates the opposite. Whether or not a structure is visible in the flow depends on the energy it captures. Near $\alpha = 0$, the output mode captures nearly as much as is possible, while at higher α , it captures substantially less. Capturing energy at low α is more important as the growth is the greatest here, so the leading output mode does a good job of predicting the structures observed in Poiseuille flow (Abe, Kawamura & Matsuo 2001). The large discontinuity in figure 11(a) in the lower right of the plot occurs because the peak time for these wavenumbers is $t = 0$, so the structures at this time only depend on the initial correlation, not on the linearized Navier–Stokes operator, so the output modes here cannot hope to capture any structure.

Figure 11(b) compares the output modes with the POD modes via the inner product, again at the peak time in G^{mean} . At each wavenumber pair, the plot shows the matrix of square inner products $G_{ij} = |\langle \mathbf{u}^i, \boldsymbol{\phi}^j \rangle|^2$ up to three modes in each basis. For $\alpha = 0$, this matrix is nearly diagonal, indicating, once again, that the output modes and POD modes are very similar, even for the subleading modes. At higher α , the modes become less similar as shown by the off-diagonal terms in the matrix of inner products.

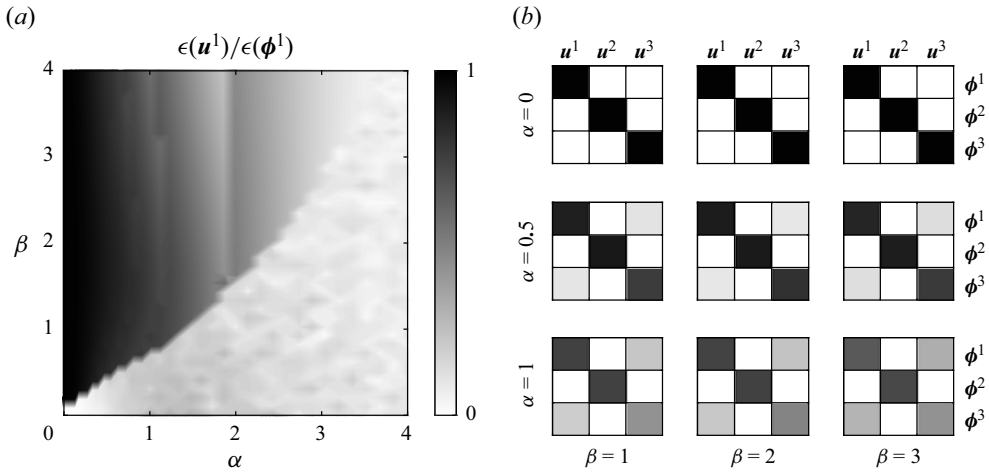


Figure 11. Comparison of the POD and output modes. (a) The ratio of energy captured by the first output mode to that of the first POD mode for a range of wavenumbers. The modes are comparable for low α . (b) The square inner products of the first three POD modes and the first three output modes. In both panels, the modes at each wavenumber are compared at the peak time in G^{mean} .

We also experimented with correlation matrices that do not respect the symmetry about $y = 0$ in the channel (arising, e.g. due to vibrations in one of the plates, but not the other). These resulted in less similarity in the modes, and the square inner products were in the range 0.5–0.9, even at $\alpha = 0$. However, so long as the correlation was symmetric, the modes were quite similar for $\alpha = 0$. The output modes are therefore a good model for the POD modes under these conditions, and leveraging the statistics may not present much advantage in predicting the structures. We stress, however, that the energy of each structure is highly dependent on the statistics, so the SVD of the evolution operator does not provide the associated energies accurately. Whether the POD modes and output modes for other flows coincide to the extent that they do in Poiseuille flow may be an interesting topic for future investigation.

It can be shown that the POD modes in flows with homogeneous directions are still delta functions in wavenumber space in those directions (Lumley 1967). Therefore, despite the fact that the disturbances will not themselves be delta functions in wavenumber, the similarity in the POD modes and output modes observed in this subsection still applies in the three-dimensional case. However, the behaviour of G^{mean} for three-dimensional disturbances can be markedly different, as shown next.

4.3. Three-dimensional disturbances: inclusion of multiple wavenumbers

Clearly, just as real initial disturbances will not identically match the maximally amplified one, real disturbances do not exist at just one pair of streamwise and spanwise wavenumbers. Parallel flow offers an analytical simplification to an analysis of transient growth – each streamwise and spanwise wavenumber pair may be considered separately in its ability to produce growth. However, this tempts further exaggeration of the possibility for large-scale transient growth. In modal stability, one need not add these wavenumbers back together to get an answer as to the long-term behaviour – if any wavenumber pair grows exponentially, so will the entire disturbance. However, if one wavenumber pair experiences large transient growth, this only implies a large gain for the entire disturbance to the extent that its initial energy is concentrated at that wavenumber pair. In this

subsection, we incorporate a range of wavenumbers and show that the weight for each pair is determined by the Fourier transform in x and z of the three-dimensional correlation. When these three-dimensional correlations are incorporated, substantially less growth is observed. We also observe a linear scaling with Reynolds number for an isotropic correlation in contrast with the quadratic scaling observed for G_{max}^{opt} and for G_{max}^{mean} at a particular α and β .

One way to calculate G^{mean} for disturbances containing multiple wavenumbers would be to define a large domain in x and z (to approximate the desired infinite directions), calculate \mathbf{A} and define a discrete correlation matrix for the full three-dimensional problem. The mean energy amplification would then be given by (4.7) and the correlation matrix by (4.15). However, this strategy is needlessly computationally intensive because it does not take advantage of the analytic simplification possible in parallel flow. Instead, we can add the expected energies at each wavenumber together, modulated by the energy of the incoming disturbances at each wavenumber. Denoting the disturbance discretized in y , but continuous in x and z as $\mathbf{q}_t(x, z)$, its energy is

$$\|\mathbf{q}(t)\|^2 = \text{Tr} \left\{ \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathbf{q}_t(x, z) \mathbf{q}_t^*(x, z) \, dx \, dz \right\} \tag{4.23}$$

$$= \frac{1}{4\pi^2} \text{Tr} \left\{ \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathbf{q}_t(\alpha, \beta) \mathbf{q}_t^*(\alpha, \beta) \, d\alpha \, d\beta \right\} \quad (\text{Parseval}) \tag{4.24}$$

$$= \frac{1}{4\pi^2} \text{Tr} \left\{ \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathbf{M}_t(\alpha, \beta) \mathbf{q}_0(\alpha, \beta) \mathbf{q}_0^*(\alpha, \beta) \mathbf{M}_t^*(\alpha, \beta) \, d\alpha \, d\beta \right\}. \tag{4.25}$$

Finally, taking an expected value, dividing by the expected initial energy, and incorporating the weight matrix yields

$$G^{mean}(t) = \frac{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \text{Tr} \{ \mathbf{L} \mathbf{M}_t(\alpha, \beta) \hat{\mathbf{C}}_{00}(\alpha, \beta) \mathbf{M}_t^*(\alpha, \beta) \mathbf{L}^* \} \, d\alpha \, d\beta}{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \text{Tr} \{ \mathbf{L} \hat{\mathbf{C}}_{00}(\alpha, \beta) \mathbf{L}^* \} \, d\alpha \, d\beta}. \tag{4.26}$$

Here, $\hat{\mathbf{C}}_{00}(\alpha, \beta) = \mathbb{E}[\mathbf{q}_0(\alpha, \beta) \mathbf{q}_0^*(\alpha, \beta)]$ is the y -discretized correlation of the initial disturbance at each wavenumber pair. By the Wiener–Khinchin theorem (Wiener 1930; Khintchine 1934), $\hat{\mathbf{C}}_{00}(\alpha, \beta)$ is, equivalently, the Fourier transform of the three-dimensional correlation

$$\hat{\mathbf{C}}_{00}(\alpha, \beta) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathbf{C}_{00}(x, z) \exp(-i(\alpha x + \beta z)) \, dx \, dz. \tag{4.27}$$

4.4. Numerical experiments using disturbances with a distribution of wavenumbers

To maximize the potential for growth, we again choose only the wall-normal-velocity autocorrelation to be non-zero.

4.4.1. Isotropic correlation

We begin by taking the wall-normal autocorrelation to be an isotropic Gaussian with correlation length λ

$$\mathbf{C}_{00}^{vv}(\mathbf{x}_1, \mathbf{x}_2) = A \exp \left[-\frac{1}{\lambda^2} (|\mathbf{x}_2 - \mathbf{x}_1|^2) \right], \tag{4.28}$$




Figure 12. Value of G_{max}^{mean} for a three-dimensional isotropic correlation with correlation length λ at $Re = 1000$. Even at the optimal correlation length, the inclusion of all wavenumbers causes G_{max}^{mean} to be roughly 2.5% of G_{max}^{opt} at $\alpha = 0, \beta = 2$.

where $|\cdot|$ denotes Euclidean distance and, again, A has no impact on G^{mean} . When discretized in y , the correlation becomes $\mathbf{C}_{00}(\Delta x, \Delta z) = \mathbf{C}_{00}^y \exp[-1/\lambda^2(\Delta x^2 + \Delta z^2)]$, where \mathbf{C}_{00}^y is the discretization of the y -dependent part of the correlation. The Fourier transform of the correlation in x and z is

$$\hat{\mathbf{C}}_{00}(\alpha, \beta) = \pi \lambda^2 A \exp\left[-\frac{\lambda^2}{4}(\alpha^2 + \beta^2)\right] \mathbf{C}_{00}^y. \tag{4.29}$$

Inserting (4.29) into (4.26) gives

$$G^{mean}(t) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \text{Tr}\{\mathbf{L}\mathbf{M}_t(\alpha, \beta)\mathbf{C}_{00}^y\mathbf{M}_t^*(\alpha, \beta)\mathbf{L}^*\} \exp\left[-\frac{\lambda^2}{4}(\alpha^2 + \beta^2)\right] d\alpha d\beta \\ \Big/ \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \text{Tr}\{\mathbf{L}\mathbf{C}_{00}^y\mathbf{L}^*\} \exp\left[-\frac{\lambda^2}{4}(\alpha^2 + \beta^2)\right] d\alpha d\beta. \tag{4.30}$$

The exponential term can be interpreted as the expected energy at each wavenumber pair implied by the correlation. For Poiseuille flow at $Re = 1000$, the most amplified wavenumbers are near $\alpha = 0, \beta = 2$ (see figures 2 and 7), and the amplification drops off rapidly as α moves away from zero. To concentrate energy near $\alpha = 0$, the correlation length must be quite long. This longer correlation also promotes growth because, as we described in the previous section, the longer the correlation length in y , the more growth is observed. However, with a long correlation length, energy is concentrated near $\beta = 0$, which does not experience much growth (see figures 2 and 7). Also detracting from G^{mean} is the fact that the maxima occur at significantly different times for different wavenumbers (see figure 6). The combined effect of these factors can be seen in figure 12. At $Re = 1000$, even with the correlation length that promotes the most growth ($\lambda^{-1} = 1.2$), G_{max}^{mean} is only 2.5% of G_{max}^{opt} at the optimal wavenumbers. The maximizing value of $\lambda^{-1} = 1.2$ can be understood as balancing energy concentration at growth-promoting values of α while also doing so at growth-promoting values of β . A small λ^{-1} leads to high energy concentration near $\alpha = 0$, which promotes growth, but also leads to low energy concentration near $\beta = 2$, which limits growth. Conversely, if $\lambda^{-1} \approx 2$, less energy is concentrated at $\alpha = 0$, which is growth limiting, but more is concentrated at $\beta = 2$, which is growth promoting.




Figure 13. Value of G_{max}^{mean} for an isotropic initial correlation for a range of correlation lengths and Reynolds numbers. Similar dependence on correlation length λ is observed at all Re , and all values are substantially lower than their single-wavenumber counterparts in figure 8.




Figure 14. Scaling of G_{max}^{mean} for the isotropic correlation for various correlation lengths along with that of G_{max}^{opt} . The grey guidelines show linear and quadratic Reynolds number scaling. Unlike G_{max}^{opt} , G_{max}^{mean} with the isotropic correlation scales linearly with Reynolds number.

The observed optimum of $\lambda^{-1} = 1.2$ strikes a balance between concentrating energy near $\alpha = 0$ and $\beta = 2$.

Figure 13 shows contours of G_{max}^{mean} for a range of correlation lengths and Reynolds numbers. The vertical dashed line divides the asymptotically stable and unstable regions. At Reynolds numbers higher than this, G_{max}^{mean} is technically infinite, but there is an initial peak in $G_{max}^{mean}(t)$ long before the instability dominates. In this figure, and all subsequent ones which show G_{max}^{mean} above the critical Reynolds number, we plot the magnitude of the initial peak in G_{max}^{mean} . The effect of the correlation length is similar across Reynolds numbers. Comparing figure 13 with figure 8 (its single-wavenumber analogue), we see that the isotropic correlation matrix severely limits growth at all Reynolds numbers. Indeed, the difference becomes greater as the Reynolds number increases.

The Reynolds number scaling is shown in figure 14. Unlike G_{max}^{opt} or G_{max}^{mean} at a particular wavenumber pair, G_{max}^{mean} for an isotropic three-dimensional correlation scales

nearly linearly with Reynolds number. This is a surprising result – the three-dimensional $G^{mean}(t)$ is obtained in (4.26) by integrating over $G^{mean}(t)$ at particular wavenumbers. However, critically, these single-wavenumber values for G^{mean} peak at different times, as can be seen in figure 6. The difference in the scaling means that the difference between G_{max}^{mean} and G_{max}^{opt} becomes larger with Reynolds number, i.e. G_{max}^{opt} increasingly overpredicts the mean energy amplification with increasing Reynolds number.

4.4.2. Non-isotropic correlation

As a generalization of the isotropic correlation investigated above, we next consider the ellipsoid

$$\mathbf{C}^{vv}(\mathbf{x}_1, \mathbf{x}_2) = A \exp \left[- \left(\frac{\Delta x^2}{\lambda_x^2} + \frac{\Delta y^2}{\lambda_y^2} + \frac{\Delta z^2}{\lambda_z^2} \right) \right], \quad (4.31)$$

where λ_x , λ_y and λ_z are the correlation lengths in the streamwise, spanwise and wall-normal directions, respectively. For example, this allows for the correlations to persist longer in x than in y or z , as may result from the advective nature of the flow (He, Jin & Yang 2017). With this extra freedom relative to the isotropic case, we may ask whether the Reynolds number scaling remains linear, as it is in that case, or becomes quadratic, as it is for G^{opt} . The answer depends on the correlation lengths chosen, but we find that if we fix λ_x at some non-zero value and vary Reynolds number, the scaling is linear.

When discretized in y , the ellipsoid correlation becomes

$$\mathbf{C}_{00}(\Delta x, \Delta z) = \mathbf{C}_{00}^y \exp \left[- \left(\frac{\Delta x^2}{\lambda_x^2} + \frac{\Delta z^2}{\lambda_z^2} \right) \right], \quad (4.32)$$

where, again, \mathbf{C}_{00}^y is the discretized y -dependent part. Upon taking the Fourier transform, the correlation in wavenumber space is

$$\hat{\mathbf{C}}_{00}(\alpha, \beta) = \pi \lambda_x \lambda_z A \exp \left[-\frac{1}{4}(\lambda_x^2 \alpha^2 + \lambda_z^2 \beta^2) \right] \mathbf{C}_{00}^y. \quad (4.33)$$

Once again, we make use of (4.26) to obtain

$$\begin{aligned} G^{mean}(t) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \text{Tr}\{\mathbf{L}\mathbf{M}_t(\alpha, \beta)\mathbf{C}_{00}^y\mathbf{M}_t^*(\alpha, \beta)\mathbf{L}^*\} \exp \left[-\frac{1}{4}(\lambda_x^2 \alpha^2 + \lambda_z^2 \beta^2) \right] d\alpha d\beta \\ &\quad / \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \text{Tr}\{\mathbf{L}\mathbf{C}_{00}^y\mathbf{L}^*\} \exp \left[-\frac{1}{4}(\lambda_x^2 \alpha^2 + \lambda_z^2 \beta^2) \right] d\alpha d\beta. \end{aligned} \quad (4.34)$$

The optimal correlation lengths and the resulting maximum value of G_{max}^{opt} are shown in figure 15. Plotted on the left axis in blue, G_{max}^{mean} scales quadratically when the correlation lengths are optimized. At $Re = 1000$, $G_{max}^{mean} = 54$, which is 27% of G_{max}^{opt} at the same Reynolds number. The optimal correlation lengths, plotted on the right axis in green, do not vary significantly with Reynolds number. The optimal λ_y is close to the channel half-height, which is consistent with the most amplified wavenumbers shown in figure 6. The maximizing λ_x^{-1} is zero. This means that the disturbances are infinitely correlated in x , implying that their energy is concentrated at $\alpha = 0$. This wavenumber is known to produce the most growth (see, e.g. figures 7 and 2), so it is not surprising that the optimal λ_x^{-1} is zero. The characteristic β in the initial correlation is $2\lambda_z^{-1}$, so in light of figure 7, the maximizing λ_z^{-1} is not a surprise. The optimal correlation lengths are not reported in




Figure 15. Value of G_{max}^{mean} optimized over the three correlation lengths (left axis) and the optimal correlation lengths (right axis) vs Re . The scaling of G_{max}^{mean} is quadratic here, and the optimal correlations do not change substantially with Re .




Figure 16. Scaling of G_{max}^{mean} with Re for $[\lambda_y^{-1}, \lambda_z^{-1}] = [1, 1.7]$ and various λ_x . The scaling is initially quadratic but becomes linear at higher Re . Longer correlations in x remain quadratic to higher Re .

figure 15 when $G_{max}^{mean} = 1$; this only occurs if $G^{mean}(t)$ peaks at the initial time $t = 0$, in which case any set of correlation lengths will produce the same result.

The correlation will, in reality, be finite in x , and this can affect the scaling. Figure 16 shows G_{max}^{mean} vs Reynolds number for constant finite choices of λ_x, λ_y and λ_z . The values of λ_x are shown and $[\lambda_y^{-1}, \lambda_z^{-1}] = [1, 1.7]$, which are near their optimal values, as shown in figure 15. For small Reynolds numbers, the scaling is nearly quadratic, matching that of the infinitely correlated disturbances, but as the Reynolds number increases, the scaling once again becomes linear. The Reynolds number at which the scaling changes depends on λ_x , with longer correlations remaining quadratic up to higher Reynolds numbers.

5. Estimating the probability density function

As figures 6 and 12 show, the expected value of the energy may be orders of magnitude smaller than G_{max}^{opt} for certain reasonable incoming correlations. With a formula for the expected energy (4.7), one may still ask whether large gains are possible or negligibly

unlikely. The linearized Navier–Stokes equations are constantly disturbed for a real flow, so the fact that G^{mean} may be small compared with G^{opt} does not preclude one of the many disturbances from achieving near-optimal growth. How long must one wait to see a dangerous disturbance? Does a short eddy turnover time in the flow imply that near-optimal growth is inevitable regardless of how small G^{mean} may be?

In order to answer questions such as these, we investigate the p.d.f. of the disturbance energy. Whereas the mean energy depends only on the initial correlations, the p.d.f. depends on the entire distribution of the incoming disturbances. We first describe a basic Monte Carlo approach for estimating the p.d.f. and apply this to two candidate distributions for the incoming disturbances, noticing that the p.d.f. of the energy drops nearly exponentially. In the two subsequent subsections, two distributions for the initial disturbances are considered, a multivariate Gaussian and a transformation of the uniform distribution on the N -sphere. For these distributions, it is possible to analytically approximate the p.d.f. of the energy using the exactly calculable moments of the energy distribution. With an accurate estimate of the p.d.f., we can calculate percentile curves for the trajectories.

We denote the p.d.f. of a random variable X as $f_X(x) \equiv \lim_{dx \rightarrow 0} Pr\{X \in [x, x + dx]\}/dx$, where $Pr\{\cdot\}$ denotes probability. If $X \in \mathbb{R}^N$ is a vector, then the p.d.f. is defined as

$$f_X(\mathbf{x}) = f_X(x_1, \dots, x_N) \equiv \lim_{dx_1 \rightarrow 0, \dots, dx_N \rightarrow 0} \frac{Pr\{X_1 \in [x_1, x_1 + dx_1], \dots, X_N \in [x_N, x_N + dx_N]\}}{dx_1 \dots dx_N}. \tag{5.1}$$

The incoming disturbances follow some distribution

$$\mathbf{q}(0) \sim f_{\mathbf{q}(0)}(\mathbf{q}_0), \tag{5.2}$$

and this implies a distribution

$$\mathbf{q}(t) \sim f_{\mathbf{q}(t)}(\mathbf{q}_t), \tag{5.3}$$

for the disturbances $\mathbf{q}(t) = \mathbf{M}_t \mathbf{q}(0)$ at time t . These distributions are the most descriptive statistical information about the disturbances; any statistic of the disturbances is implied by the full distribution. For example, the correlation matrix \mathbf{C}_{00} for the initial disturbances is implied by the distribution of the initial disturbance $f_{\mathbf{q}(0)}$. The converse, however, is not true – there are many distributions with the same correlation matrix. In the last section, we showed that G^{mean} only depends on the correlation matrix, so there was no need to consider the form of the underlying distribution. However, to calculate the p.d.f. of the energy of the disturbance at some time,

$$e(t) \sim f_{e(t)}(e_t), \tag{5.4}$$

the full distribution of initial disturbances is needed.

5.1. Monte Carlo

With a means of sampling initial disturbances from $f_{\mathbf{q}(0)}(\mathbf{q}_0)$, samples of the growing disturbances can be generated by multiplying the initial ones by \mathbf{M}_t , and samples of $e(t)$ are finally obtained by computing their norm. An estimation of the p.d.f. can be obtained using standard methods, such as `ksdensity` in MATLAB. Figure 17 shows the empirical p.d.f. resulting from performing this Monte Carlo with two different distributions of the initial disturbances (described later), both with the same correlation. The Monte Carlo is performed at $t = 100$, $Re = 1000$, $\alpha = 0$, $\beta = 2$ with a correlation of the form (4.22) and correlation length $\lambda^{-1} = 5$. The distributions are described in detail in the




Figure 17. Empirical p.d.f. of energy at $t = 100$ resulting from initial disturbances distributed as a transformation of the uniform distribution and a multivariate Gaussian, both with the same correlation with correlation length $\lambda^{-1} = 5$. Note the similarity in the two p.d.f.s and the near-exponential decay.

following subsections. Two observations are apparent. First, both distributions result in a very similar p.d.f. for the energy of the amplified disturbance. That the p.d.f.s are similar indicates that while the p.d.f. of the energy is a function of the full distribution of incoming disturbances, reasonable distributions of initial disturbances with the same correlation will give similar p.d.f.s for the energy. Second, the p.d.f.s decay nearly exponentially. The exponential decay indicates that it is very unlikely that the energy of an amplified disturbance substantially exceeds $G^{mean}(t)$. The exponential decay also allows for accurate *a priori* approximation of the p.d.f., which we discuss in the following subsections for the two distributions of initial disturbances.

5.2. Multivariate Gaussian

First, we assume the initial disturbances follow the multivariate Gaussian with mean $\mathbf{0}$ and correlation \mathbf{C}_{00}

$$\mathbf{q}(0) \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_{00}). \tag{5.5}$$

Also, we assume the correlation to have unit trace, so $G^{mean}(t) = \mathbb{E}[e(\mathbf{q}(t))]$. More explicitly, this distribution, as represented by its p.d.f., is

$$f_{\mathbf{q}(0)}(\mathbf{q}_0) = \frac{\exp\left[-\frac{1}{2}\mathbf{q}_0^* \mathbf{C}_{00}^{-1} \mathbf{q}_0\right]}{\sqrt{(2\pi)^N |\mathbf{C}_{00}|}}. \tag{5.6}$$

If \mathbf{C}_{00} is rank deficient, the inverse \mathbf{C}_{00}^{-1} and determinant $|\mathbf{C}_{00}|$ are modified to the pseudoinverse and pseudodeterminant.

Any linear function of $\mathbf{q}(0)$ also follows a multivariate Gaussian distribution (Tong 2011), so the disturbance $\mathbf{q}(t) = \mathbf{M}_t \mathbf{q}(0)$ some time in the future is distributed as

$$\mathbf{q}(t) \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_{tt}), \tag{5.7}$$

where the correlation $\mathbf{C}_{tt} = \mathbf{M}_t \mathbf{C}_{00} \mathbf{M}_t^*$ comes from (4.15). The energy is $e(t) = \|\mathbf{q}(t)\|^2$ and we seek to estimate its p.d.f. $f_{e(t)}(e_t)$. This distribution is one of a well-studied class – quadratic forms in multivariate Gaussian variables. The moments of these

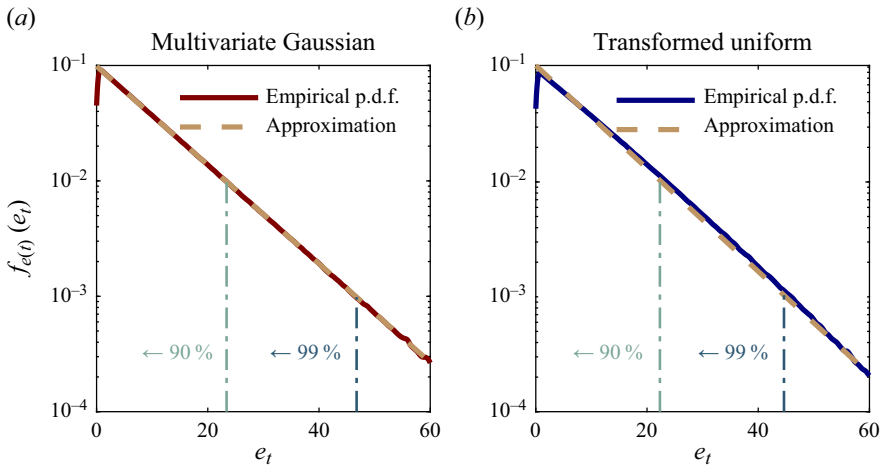


Figure 18. Empirical p.d.f. for the disturbance energy and the approximation thereof. The approximation is the exponential distribution with the same fourth moment as the true distribution and is calculated without the Monte Carlo.

distributions are known (Mathai & Provost 1992), and for the case at hand, the r th moment μ_r can be calculated recursively as

$$\mu_r \equiv \mathbb{E}[e^r] = \sum_{k=0}^{r-1} \binom{r-1}{k} g^{(r-1-k)} \mu_k, \tag{5.8a}$$

where

$$g^{(k)} = 2^k k! \text{Tr}\{(\mathbf{L}\mathbf{C}_t\mathbf{L}^*)^{k+1}\}, \tag{5.8b}$$

and $\mu_0 = 1$. Note that $\mu_1 = G^{mean}$, i.e. the first moment recovers the expected energy.

Our goal is to estimate the right tail of the p.d.f. of the energy in order to approximate the probability of exceeding a particular energy. As shown in figure 17, the right tail of the empirical p.d.f. displays nearly exponential decay. Therefore we assume its form to be

$$f_{e(t)}(e_t) \approx \gamma \exp[-\gamma e_t]. \tag{5.9}$$

To find the decay rate, we find γ such that the r th moment of the exponential ansatz matches the r th moment of the true distribution, given in (5.8). The true distribution, estimated via the Monte Carlo, is near exponential for high energies but not low ones, so to find the correct exponential parameter we equate a relatively high moment, since this weights the high-energy tail of the distribution heavily. Denoting the moment equated as r , the r th moment of the exponential distribution (5.9) is $\mu_r = r!/\gamma^r$. Equating this to the true r th moment given in (5.8) and solving for the exponential decay rate gives

$$\gamma = \left(\frac{r!}{\mu^r}\right)^{1/r}. \tag{5.10}$$

This is an analytical approximation; the only role of the previous Monte Carlo was to suggest the exponential form (5.9). Figure 18(a) shows this approximation strategy using $r = 4$. The confidence bounds are derived by integrating the approximate p.d.f. The Reynolds number, wavenumbers and correlation length are $Re = 1000$, $\alpha = 0$, $\beta = 2$, $\lambda^{-1} = 5$. Out of the 10^7 trajectories used to generate the empirical distribution, 99.005 % were below the 99 % confidence bound.

5.3. Transformation of a uniform distribution on the N -sphere

Second, we assume that the disturbances are distributed as some transformation of a uniform distribution on the surface of the N -sphere

$$\mathbf{q}(0) \sim \Psi \mathbf{u}, \tag{5.11}$$

where $\|\mathbf{u}\| = 1$ with probability 1, and $\mathbf{u} \sim \mathbf{R}\mathbf{u}$ for any rotation matrix \mathbf{R} . Samples from this distribution can be easily generated by normalizing samples from an independent and identically distributed multivariate Gaussian to be the same radius R , $\mathbf{u} = R(\mathbf{x}/\|\mathbf{x}\|)$ with $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. The transformation Ψ can be chosen so that the initial disturbances have some desired correlation \mathbf{C}_{00} . By choosing Ψ such that

$$\Psi \Psi^* = N\mathbf{C}_{00}, \tag{5.12}$$

e.g. setting it to the Cholesky decomposition of $N\mathbf{C}_{00}$, where N is the dimension of the state, the correlation matrix of the initial disturbances is

$$\mathbb{E}[\Psi \mathbf{u} \mathbf{u}^* \Psi^*] = \frac{1}{N} \Psi \Psi^* = \mathbf{C}_{00}. \tag{5.13}$$

The first equality holds because $\mathbb{E}[\mathbf{u} \mathbf{u}^*] = (1/N)\mathbf{I}$. The disturbance energy some time later is given by $e(t) = \|\mathbf{q}(t)\|^2 = \|\mathbf{T}\mathbf{u}\|^2$. The final equality expresses the energy as the square norm of the uniform distribution acted on by a matrix $\mathbf{T} = \mathbf{M}_t \Psi$.

The moments for this transformation of the uniform distribution were derived by von Neumann (1941). The result is (Kargin 2010)

$$\mathbb{E}[e^k] = \frac{2^k k!}{N(N+2)(N+4)\cdots(N+2k-2)} \xi_k, \tag{5.14}$$

where the ξ are defined as the power series coefficients of

$$1 + \xi_1 z + \xi_2 z^2 + \xi_3 z^3 + \cdots = \exp[\zeta_1 z + \zeta_2 z^2 + \zeta_3 z^3 + \cdots], \tag{5.15}$$

and ζ_j is defined in terms of the singular values $\sigma_i(\mathbf{L}\mathbf{T})$, as

$$\zeta_j = \frac{1}{2j} \sum_{i=1}^N \sigma_i^{2j}. \tag{5.16}$$

Using these moments, we can approximate the true p.d.f. using the technique described in the previous subsection. Figure 18 shows the result along with confidence bounds derived by integrating the approximated p.d.f. Just as before, the Reynolds number, wavenumbers and correlation length are $Re = 1000$, $\alpha = 0$, $\beta = 2$, $\lambda^{-1} = 5$. Of the 10^7 trajectories used to calculate the empirical distribution, 98.981 % were less than the 99 % confidence bound.

6. Conclusions

Standard transient growth analyses are based on the maximum growth experienced by any initial disturbance. While this is a useful upper bound on linear growth, it can vastly overpredict the growth experienced by real disturbances. We have developed a statistical framework to explore the space of real disturbances and quantify their growth. We demonstrated the framework and its ability to extract insights on Poiseuille flow.

The framework can be summarized using figure 19. Here, G^{opt} , the quantity used in the literature to quantify transient growth, far overshoots all one-thousand random trajectories,

the context of temporal stability, this correlation matrix is a different quantity from the correlation one might measure by time averaging in an experiment. The figures reported, e.g. for the ratio of G^{mean}/G^{opt} , are not meant to be taken as quantitative predictions of what would be observed in an experiment. Rather, they are meant to show trends and to emphasize that, when various factors are accounted for, the growth of real disturbances can be significantly smaller than G^{opt} . Although G^{mean} depends on the full correlation matrix, the correlation length is a particularly important feature in determining G^{mean} . The observation that long correlations lead to more growth may serve as a practical guide when an accurate model for correlations is unavailable.

We have also described a statistical approach to determining the structures that emerge, in the form of modes from a particular POD problem. These structures also depend on the initial correlations but were observed only to differ slightly from the standard output modes of the evolution operator, which do not depend on the correlations. This indicates that the output modes do a good job of capturing the energy of the growing disturbances regardless of the correlation matrix. Nevertheless, they do not accurately predict the energy of each structure, whereas the POD eigenvalues do.

We have discussed the statistical framework in the context of temporal stability, wherein an initial disturbance at a particular time is assumed, then evolved forward in time without further forcing to the linear dynamics. In spatial stability, a disturbance is introduced at a particular streamwise location, and its growth is then calculated as it evolves downstream as a function of the streamwise coordinate. Transient growth has been investigated in the context of spatial stability (Hack & Moin 2017; Fava *et al.* 2022), and the framework developed in this paper applies equally to spatial stability by exchanging t for x and the linearized Navier–Stokes operator for a spatial evolution operator (Towne & Colonius 2015). Spatial stability may be an excellent application for two reasons. First, applied to spatial stability, $G^{mean}(x_1)$ represents the ratio of turbulent intensities at $x = x_1$ and $x = 0$. Consider a laminar boundary layer excited at $x = 0$ by free-stream turbulence. The turbulent intensity at various points along the boundary layer is critical for determining where transition occurs, and its evolution in space is given by $G^{mean}(x)$. In this context, we have shown that the turbulent intensity depends not only on the input turbulent intensity, encapsulated by the diagonal terms of the initial correlation matrix, but on the off-diagonal terms as well. Second, the correlation matrix of the initial disturbances may be easier to model physically in spatial stability than in temporal stability. In this case, incoming turbulence is often the source of the disturbances and turbulent correlations have been studied thoroughly. In the example of a boundary layer, the free-stream turbulence that excites the leading edge may be modelled by the von Kármán spectrum (von Kármán 1948), which implies the correlations within the initial disturbances. Also, with such a model for the initial correlation, the POD modes of the space-evolved correlation may provide a better basis for the structures that arise downstream than the output modes of the matrix exponential.

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

Author ORCIDs.

 Peter Frame <https://orcid.org/0000-0002-9919-1796>;

 Aaron Towne <https://orcid.org/0000-0002-7315-5375>.

Appendix A. Stochastic formulation and connection to resolvent analysis

The framework described in this paper operates within the physical model of transient growth, wherein the linearized Navier–Stokes equations are impulsively perturbed by q_0

at $t = 0$ and are then left undisturbed. One may instead assess the importance of the non-normality by assuming a stochastic forcing to the linearized Navier–Stokes equations, rather than an impulsive one. The equations governing the disturbance are then

$$\dot{\mathbf{q}}(t) = \mathbf{A}\mathbf{q}(t) + \mathbf{f}(t), \tag{A1}$$

where $\mathbf{f}(t)$ represents the stochastic forcing, which may come from, for example, free-stream turbulence or vibrations in the channel. The degree to which non-normality amplifies the forcing is captured by the ratio of expected disturbance energy to expected forcing power

$$\Gamma^{mean} = \frac{\mathbb{E}[\|\mathbf{q}(t)\|^2]}{\mathbb{E}[\|\mathbf{f}(t)\|^2]}. \tag{A2}$$

To calculate Γ^{mean} , it is useful to work in the frequency domain. It is well known that the Fourier-transformed disturbance is then related to the Fourier-transformed forcing by

$$\hat{\mathbf{q}}(\omega) = \mathbf{R}(\omega)\hat{\mathbf{f}}(\omega), \tag{A3}$$

where $\mathbf{R}(\omega) = (i\omega\mathbf{I} - \mathbf{A})^{-1}$ is the resolvent. This equation has been studied extensively (Trefethen *et al.* 1993; Farrell & Ioannou 1996; McKeon & Sharma 2010; McKeon 2017). In the context of quantifying the non-normality, the quantities usually reported are the first few singular values of $\mathbf{R}(\omega)$, often maximized over frequency. The associated left singular vector can indeed be remarkably useful for predicting structures that appear in the flow (Sharma & McKeon 2013; Luhar, Sharma & McKeon 2014; Schmidt *et al.* 2018), but entirely analogous to the optimal transient growth case, the singular values of the resolvent operator overstate the degree to which real forcing is amplified. This occurs because the real forcing will not exactly coincide with the right singular vector of the resolvent operator, and it will not oscillate at exactly the optimal frequency.

By multiplying (A3) by its conjugate transpose and taking the expected value, the cross-spectral density of the disturbance $\mathbf{S}_{qq}(\omega) = \mathbb{E}[\hat{\mathbf{q}}(\omega)\hat{\mathbf{q}}^*(\omega)]$ may be related to that of the forcing $\mathbf{S}_{ff}(\omega) = \mathbb{E}[\hat{\mathbf{f}}(\omega)\hat{\mathbf{f}}^*(\omega)]$ in terms of the resolvent as (Towne *et al.* 2018)

$$\mathbf{S}_{qq}(\omega) = \mathbf{R}(\omega)\mathbf{S}_{ff}(\omega)\mathbf{R}^*(\omega). \tag{A4}$$

This is useful for obtaining the expected disturbance energy because, as a consequence of the Wiener–Khinchin theorem, the disturbance autocorrelation at zero time lag, $\mathbf{C}_{qq}(0) = \mathbb{E}[\mathbf{q}(t)\mathbf{q}^*(t + 0)]$, may be related to the disturbance cross-spectral density as

$$\mathbf{C}_{qq}(0) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \mathbf{S}_{qq}(\omega) \, d\omega. \tag{A5}$$

The (weighted) trace of this correlation is the expected disturbance energy, $\mathbb{E}[\|\mathbf{q}(t)\|^2] = \text{Tr}\{\mathbf{L}\mathbf{C}_{qq}(0)\mathbf{L}^*\}$. Combining (A4) and (A5), Γ^{mean} may be written as

$$\Gamma^{mean} = \frac{\int_{-\infty}^{\infty} \text{Tr}\{\mathbf{L}\mathbf{R}(\omega)\mathbf{S}_{ff}(\omega)\mathbf{R}^*(\omega)\mathbf{L}^*\} \, d\omega}{\int_{-\infty}^{\infty} \text{Tr}\{\mathbf{L}\mathbf{S}_{ff}(\omega)\mathbf{L}^*\} \, d\omega}. \tag{A6}$$

Modelling the forcing cross-spectral density $\mathbf{S}_{ff}(\omega)$ is equivalent to modelling the space–time forcing correlation $\mathbf{C}_{ff}(\tau) = \mathbb{E}[\mathbf{f}(t)\mathbf{f}^*(t + \tau)]$. Thus, in order to quantify the expected state energy, one must have knowledge of these spatio-temporal forcing statistics (Towne, Brès & Lele 2017; Morra *et al.* 2020; Karban *et al.* 2023).

Just as G^{mean} optimized over all initial disturbance correlations is G^{opt} , the maximum of Γ^{mean} over all $\mathbf{S}_{ff}(\omega)$ is the square of the first singular value of the weighted resolvent, maximized over all frequencies, $\sigma_1^2(\omega_{max}) = \max_{\omega} \sigma_1^2(\omega)$. This occurs when the cross-spectral density is (proportional to) the outer product of the right singular vector of the weighted resolvent operator at that frequency, $\mathbf{S}_{ff}(\omega) = \delta(\omega - \omega_{max}) \mathbf{v}_1(\omega_{max}) \mathbf{v}_1^*(\omega_{max})$.

The white-noise model proposed in Farrell & Ioannou (1993, 1994, 1996) and later used by Fontane *et al.* (2008) may be recovered as the white-noise case of the stochastic formulation presented above, i.e. where \mathbf{S}_{ff} is independent of ω . These works derive $\mathbb{E}[\|\mathbf{q}(t)\|^2]$, i.e. $\Gamma^{mean} \mathbb{E}[\|\mathbf{f}(t)\|^2]$, as the trace of the solution of a particular Lyapunov equation, and it may also be derived using the numerator of (A6) with constant \mathbf{S}_{ff} . One negative consequence of the white-noise model is that the forcing has infinite power, i.e. $\mathbb{E}[\|\mathbf{f}(t)\|^2] = \infty$, but the expected energy of the disturbance remains finite. In the same way that the spatial correlations are essential to calculating the degree of expected transient growth, the spatio-temporal correlations are needed in calculating the degree of stochastic forcing amplification.

Appendix B. Probability density function separability

If the distribution of the initial disturbance is separable in radius and direction

$$\mathbf{q}(0) \sim f_{\mathbf{q}(0)}(\mathbf{q}_0) = R(\|\mathbf{q}_0\|) \Theta\left(\frac{\mathbf{q}_0}{\|\mathbf{q}_0\|}\right), \tag{B1}$$

the ratio of expected energies G^{mean} is equal to the expected ratio of energies

$$\frac{\mathbb{E}[e(\mathbf{q}(t))]}{\mathbb{E}[e(\mathbf{q}(0))]} = \mathbb{E}\left[\frac{e(\mathbf{q}(t))}{e(\mathbf{q}(0))}\right]. \tag{B2}$$

This can be shown as follows. Assume, without loss of generality, that R and Θ are scaled such that they both integrate to unity (with the appropriate measure). Defining $r^2 \equiv \mathbf{q}_0^* \mathbf{W} \mathbf{q}_0 = \|\mathbf{L} \mathbf{q}_0\|^2$ and $\boldsymbol{\theta} \equiv \mathbf{q}_0/r$, we write the left-hand side of (B2) in terms of the distribution

$$\begin{aligned} & \frac{\mathbb{E}[e(\mathbf{q}(t))]}{\mathbb{E}[e(\mathbf{q}(0))]} \\ &= \int_0^\infty dr \int_{\|\boldsymbol{\theta}\|=1} R(r) \Theta(\boldsymbol{\theta}) \|\mathbf{L} \mathbf{M}_t r \boldsymbol{\theta}\|^2 d\sigma(\boldsymbol{\theta}) \bigg/ \int_0^\infty dr \int_{\|\boldsymbol{\theta}\|=1} R(r) \Theta(\boldsymbol{\theta}) r^2 d\sigma(\boldsymbol{\theta}) \end{aligned} \tag{B3}$$

$$= \int_0^\infty R(r) r^2 dr \int_{\|\boldsymbol{\theta}\|=1} \Theta(\boldsymbol{\theta}) \|\mathbf{L} \mathbf{M}_t \boldsymbol{\theta}\|^2 d\sigma(\boldsymbol{\theta}) \bigg/ \int_0^\infty R(r) r^2 dr \int_{\|\boldsymbol{\theta}\|=1} \Theta(\boldsymbol{\theta}) d\sigma(\boldsymbol{\theta}) \tag{B4}$$

$$= \int_{\|\boldsymbol{\theta}\|=1} \Theta(\boldsymbol{\theta}) \|\mathbf{L} \mathbf{M}_t \boldsymbol{\theta}\|^2 d\sigma(\boldsymbol{\theta}). \tag{B5}$$

Above, $\sigma(\boldsymbol{\theta})$ is the spherical measure. The right-hand side of (B2) can be written

$$\mathbb{E} \left[\frac{e(\mathbf{q}(t))}{e(\mathbf{q}(0))} \right] = \int_0^\infty dr \int_{\|\boldsymbol{\theta}\|=1} R(r) \Theta(\boldsymbol{\theta}) \frac{\|\mathbf{L}\mathbf{M}_t \mathbf{r} \boldsymbol{\theta}\|^2}{r^2} d\sigma(\boldsymbol{\theta}) \quad (\text{B6})$$

$$= \int_0^\infty R(r) dr \int_{\|\boldsymbol{\theta}\|=1} \Theta(\boldsymbol{\theta}) \|\mathbf{L}\mathbf{M}_t \boldsymbol{\theta}\|^2 d\sigma(\boldsymbol{\theta}) \quad (\text{B7})$$

$$= \int_{\|\boldsymbol{\theta}\|=1} \Theta(\boldsymbol{\theta}) \|\mathbf{L}\mathbf{M}_t \boldsymbol{\theta}\|^2 d\sigma(\boldsymbol{\theta}), \quad (\text{B8})$$

so the two are equal and (B2) holds.

REFERENCES

- ABE, H., KAWAMURA, H. & MATSUO, Y. 2001 Direct numerical simulation of a fully developed turbulent channel flow with respect to the Reynolds number dependence. *Trans. ASME J. Fluids Engng* **123** (2), 382–393.
- BUTLER, K.M. & FARRELL, B.F. 1992 Three-dimensional optimal perturbations in viscous shear flow. *Phys. Fluids* **4** (8), 1637–1650.
- CHERUBINI, S., DE PALMA, P. & ROBINET, J.-CH. 2015 Nonlinear optimals in the asymptotic suction boundary layer: transition thresholds and symmetry breaking. *Phys. Fluids* **27** (3), 034108.
- FARRELL, B.F. & IOANNOU, P.J. 1993 Stochastic forcing of the linearized Navier–Stokes equation. *Phys. Fluids A* **5** (11), 2600–2609.
- FARRELL, B.F. & IOANNOU, P.J. 1994 Variance maintained by stochastic forcing of non-normal dynamical systems associated with linearly stable shear flow. *Phys. Rev. Lett.* **72** (8), 1188–1191.
- FARRELL, B.F. & IOANNOU, P.J. 1996 *Generalized Stability Theory Part I: Autonomous Operators*, 53rd edn, pp. 2025–2040. J. Atmos. Sci.
- FAVA, T.C.L., LOBO, B.A., SCHAFFARCZYK, A.P., BREUER, M., HANIFI, A. & HENNINGSON, D. 2022 On the stability and transition to turbulence of the flow over a wind-turbine airfoil under varying free-stream turbulence intensity. In *12th International Symposium on Turbulence and Shear Flow Phenomena*. TSFP.
- FONTANE, J., BRANCHER, P. & FABRE, D. 2008 Stochastic forcing of the Lamb–Oseen vortex. *J. Fluid Mech.* **613**, 233–254.
- FRAME, P. & TOWNE, A. 2023 Space-time POD and the Hankel matrix. *PLoS ONE* **18** (8), e0289637.
- GUSTAVSSON, L.H. 1986 Excitation of direct resonances in plane Poiseuille flow. *Stud. Appl. Maths* **75** (3), 227–248.
- GUSTAVSSON, L.H. 1991 Energy growth of three-dimensional disturbances in plane Poiseuille flow. *J. Fluid Mech.* **224**, 241–260.
- HACK, M.J.P. & MOIN, P. 2017 Algebraic disturbance growth by interaction of Orr and lift-up mechanism. *J. Fluid Mech.* **829**, 112–126.
- HANIFI, A., SCHMID, P.J. & HENNINGSON, D.S. 1996 Transient growth in compressible boundary layer flow. *Phys. Fluids* **8** (3), 826–837.
- HE, G., JIN, G. & YANG, Y. 2017 Space-time correlations and dynamic coupling in turbulent flow. *Annu. Rev. Fluid Mech.* **49** (1), 51–70.
- HERBERT, T. 1977 Die neutrale fläche der ebenen Poiseuille–Strömung. habilitation, Uni Stuttgart.
- KARBAN, U., BUGEAT, B., TOWNE, A., LESSHAFFT, L., AGARWAL, A. & JORDAN, P. 2023 An empirical model of noise sources in subsonic jet. *J. Fluid Mech.* **965**, A18.
- KARGIN, V. 2010 Products of random matrices: dimension and growth in norm. *Ann. Appl. Probab.* **20** (3), 890–906.
- VON KÁRMÁN, T. 1948 Progress in the statistical theory of turbulence. *Proc. Natl Acad. Sci. USA* **34** (11), 530–539.
- KERSWELL, R.R. 2018 Nonlinear nonmodal stability theory. *Annu. Rev. Fluid Mech.* **50** (1), 319–345.
- KHINTCHINE, A. 1934 Korrelationstheorie der stationären stochastischen prozesse. *Math. Ann.* **109**, 604–615.
- LUHAR, M., SHARMA, A.S. & MCKEON, B.J. 2014 On the structure and origin of pressure fluctuations in wall turbulence: predictions based on the resolvent analysis. *J. Fluid Mech.* **751**, 38–70.
- LUMLEY, J.L. 1967 The structure of inhomogeneous turbulent flow. In *Atmospheric Turbulence and Radio Wave Propagation* (ed. A.M. Yaglom & V.I. Tatarski), pp. 166–178. Nauka.
- LUMLEY, J.L. 1970 *Stochastic tools in turbulence*. Academic Press.

- MALKUS, W.V.R. 1956 Outline of a theory of turbulent shear flow. *J. Fluid Mech.* **1** (5), 521–539.
- MARKEVICIUTE, V. 2022 Statistical stability and fast transient growth in wall-bounded turbulence. PhD thesis, University of Cambridge.
- MATHAI, A. & PROVOST, S. 1992 Quadratic forms in random variables: Theory and Applications. Marcel Dekker.
- MCKEON, B.J. 2017 The engine behind (wall) turbulence: perspectives on scale interaction. *J. Fluid Mech.* **817**, P1.
- MCKEON, B.J. & SHARMA, A.S. 2010 A critical-layer framework for turbulent pipe flow. *J. Fluid Mech.* **658**, 336–382.
- MORRA, P., NOGUEIRA, P.A.S., CAVALIERI, A.V.G. & HENNINGSON, D.S. 2020 The colour of forcing statistics in resolvent analyses of turbulent channel flow. *J. Fluid Mech.* **907**, A24.
- VON NEUMANN, J. 1941 Distribution of the ratio of the mean square successive difference to the variance. *Ann. Math. Statist.* **12** (4), 367–395.
- 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.
- PRINGLE, C.C.T. & KERSWELL, R.R. 2010 Using nonlinear transient growth to construct the minimal seed for shear flow turbulence. *Phys. Rev. Lett.* **105**, 154502.
- REDDY, S.C. & HENNINGSON, D.S. 1993 Energy growth in viscous channel flow. *J. Fluid Mech.* **252**, 209–238.
- REYNOLDS, O. 1883 An experimental investigation of the circumstances which determine whether the motion of water shall be direct or sinuous, and of the law of resistance in parallel channel. *Phil. Trans. R. Soc. Lond.* **174**, 935–982.
- ROWLEY, C.W. & DAWSON, S.T.M. 2017 Model reduction for flow analysis and control. *Annu. Rev. Fluid Mech.* **49** (1), 387–417.
- SCHMID, P. & HENNINGSON, D. 2001 *Stability and Transition in Shear Flows*. Springer.
- SCHMID, P.J. 2007 Nonmodal stability theory. *Annu. Rev. Fluid Mech.* **39** (1), 129–162.
- SCHMID, P.J. & HENNINGSON, D.S. 1994 Optimal energy density growth in Hagen–Poiseuille flow. *J. Fluid Mech.* **277**, 197–225.
- SCHMIDT, O.T., TOWNE, A., RIGAS, G., COLONIUS, T. & BRÈS, G.A. 2018 Spectral analysis of jet turbulence. *J. Fluid Mech.* **855**, 953–982.
- SHARMA, A.S. & MCKEON, B.J. 2013 On coherent structure in wall turbulence. *J. Fluid Mech.* **728**, 196–238.
- SIROVICH, L. 1987 Turbulence and the dynamics of coherent structures. I – coherent structures. II – symmetries and transformations. III – dynamics and scaling. *Q. Appl. Maths* **45**, 561–571.
- 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.
- TILLMARK, N. & ALFREDSSON, P.H. 1992 Experiments on transition in plane Couette flow. *J. Fluid Mech.* **235**, 89–102.
- TONG, Y.L. 2011 *The Multivariate Normal Distribution*. Springer.
- TOWNE, A., BRÈS, G.A. & LELE, S.K. 2017 A statistical jet-noise model based on the resolvent framework. In *23rd AIAA/CEAS Aeroacoustics Conference*. American Institute of Aeronautics and Astronautics.
- TOWNE, A. & COLONIUS, T. 2015 One-way spatial integration of hyperbolic equation. *J. Comput. Phys.* **300**, 844–861.
- TOWNE, A., SCHMIDT, O.T. & COLONIUS, T. 2018 Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis. *J. Fluid Mech.* **847**, 821–867.
- TREFETHEN, L.N. & EMBREE, M. 2005 *Spectra and Pseudospectra*. Princeton University Press.
- TREFETHEN, L.N., TREFETHEN, A.E., REDDY, S.C. & DRISCOLL, T.A. 1993 Hydrodynamic stability without eigenvalue. *Science* **261** (5121), 578–584.
- WIENER, N. 1930 Generalized harmonic analysis. *Acta Mathematica* **55**, 117–258.