Adjoint chaos via cumulant truncation

John Craske[1]

[1]Department of Civil and Environmental Engineering, Imperial College London,
London SW7 2AZ, UK

(Last updated: April 3, 2019)

We describe a simple and systematic method for obtaining approximate sensitivity information from a chaotic dynamical system using a hierarchy of cumulant equations. The resulting forward and adjoint systems yield information about gradients of functionals of the system and do not suffer from the convergence issues that are associated with the tangent linear representation of chaotic systems. The functionals on which we focus are ensemble-averaged quantities, whose dynamics are not necessarily chaotic; hence we analyse the system's statistical state dynamics, rather than individual trajectories. The approach is designed for extracting parameter sensitivity information from the detailed statistics that can be obtained from direct numerical simulation or experiments. We advocate a data-driven approach that incorporates observations of a system's cumulants to determine an optimal closure for a hierarchy of cumulants that does not require the specification of model parameters. Whilst the sensitivity information from the resulting surrogate model is approximate, the approach is designed to be used in the analysis of turbulence, whose number of degrees of freedom and complexity currently prohibits the use of more accurate techniques. Here we apply the method to obtain functional gradients from low-dimensional representations of Rayleigh-Bénard convection.

1 Introduction

Complete information about a particular solution of an engineering problem is often less useful than knowledge of the way in which a small number of functionals of the solution change with respect to input parameters. An example in fluid mechanics is the effect that a body's shape has on the drag to which it is subjected (Pironneau, 1974; Jameson, 1988). Further examples can be found in the fields of data assimilation (Dimet & Talagrand, 1986), uncertainty quantification (Cacuci, 2003), stability analysis (Luchini & Bottaro, 2014; Farrell *et al.*, 2014), flow reconstruction (Foures *et al.*, 2014) and flow optimisation more generally (Lions, 1971). In these situations it is natural to focus on adjoint variables, which represent the derivative of a given functional with respect to the problem's constraints or governing equations. With adjoint variables the derivative of the functional with respect to any combination of input parameters can be readily computed with a single dot product, alleviating the need to run a large ensemble of simulations to obtain gradients in different directions. For a general introduction to the theory the reader is referred to Marchuk (1995) and Giles & Pierce (2000).

Whilst adjoint analysis is well established and used successfully in many fields, the problem of obtaining functional gradients from chaotic dynamical systems, such as turbulence, is an open question (Vishnampet *et al.*, 2015). Whether such gradients are well-defined depends on the properties of the dynamical system. For example, if the system is uniformly hyperbolic (Smale, 1967; Eckmann & Ruelle, 2004) then linear response theory provides the required formula (Eyink *et al.*, 2004; Ruelle, 2009). In all chaotic systems, however, the linearised description, on which both forward and adjoint analysis is based, produces divergent trajectories that make it impossible to compute accurate gradients over large times in the conventional way (Lea *et al.*, 2000). A variety of different

methods have been proposed to overcome this practical difficulty. A possible approximation is to obtain an estimation of a system's linear response by taking finite differences (Russo & Luchini, 2016). An approach employing adjoint formalism is to limit the duration over which sensitivity information is obtained (Vishnampet *et al.*, 2015), or to collect an ensemble of gradients to compensate for the short time intervals to which the adjoint equations are otherwise restricted (Lea *et al.*, 2000; Eyink *et al.*, 2004). In addition to the requirement of having to obtain a potentially large ensemble, the difficulty of the latter approach is in determining an appropriate time interval *a priori*. Consequently, probability density functions have also received attention as a reliable source of gradient information. Thuburn (2005) proposed solving an adjoint Fokker-Planck equation, which, though capable of producing accurate derivatives, is computationally expensive and involves approximation in the selection of stochastic forcing terms. Related work uses ideas from the fluctuation-dissipation theorem (Marconi *et al.*, 2008) to determine sensitivities (e.g. Cooper & Haynes, 2011), and typically relies on an assumption about the underlying probability density function.

Recent efforts to reconcile adjoint techniques and chaotic systems have focused on deriving sensitivities from shadow trajectories, which are defined as remaining uniformly close to a given trajectory of the system over time (Wang, 2013), and can therefore yield meaningful sensitivity information. An improvement of the method proposed by Wang (2013), which relied on the calculation of Lyapunov exponents and was therefore restricted to low-dimensional dynamical systems, is the least-squares shadowing method proposed by Wang *et al.* (2014). The least-squares shadowing method involves solving an optimisation problem to determine a perturbed trajectory that is closest to the chosen reference trajectory. Notably, the least-squares shadowing method has been applied to the Kuramoto-Sivashinsky equation, and yields accurate gradient information for certain states (Blonigan & Wang, 2014).

The issue regarding divergent trajectories in tangent and adjoint systems can be circumvented altogether by computing sensitivities of unstable periodic orbits (Lasagna, 2018). Perturbations of unstable periodic orbits, which behave like a skeleton around chaotic orbits (see e.g. Auerbach *et al.*, 1987), provide a proxy for the latter's sensitivity. In general, each unstable period orbit returns a different sensitivity. In certain cases, however, the sensitivities are closely aligned and exhibit a good agreement with that of the underlying chaotic orbit (Lasagna, 2018). Principal among the challenges associated with this technique is the difficulty of finding unstable periodic orbits in chaotic systems of high dimension, such as turbulence at high Reynolds number (see e.g. Lucas & Caulfield, 2017).

The need to overcome the sensitive dependence on perturbations inherent in chaotic systems might be regarded as unnecessary, in view of the fact that one is often interested in gradients of *ensemble-averaged* quantities. Indeed, following Hopf (1952) and Lorenz (1967), it is possible to directly simulate a system's statistics or cumulant dynamics. With the use of the original governing equations, the cumulant equations can be derived from a single flow functional (Hopf, 1952) and provide a direct means of understanding the behaviour of a flow's statistics. In addition to their evolution being slower and not necessarily chaotic, the cumulant equations can be used to investigate statistically unsteady problems, statistical stability and to provide an analytical means of determining the linear response of a system (Farrell & Ioannou, 2014). The evolution of a finite set of dependent variables

corresponds to an infinite hierarchy of cumulant equations. The benefits of focusing on the evolution of statistics are therefore offset by the requirement of finding a suitable closure (Rothmayer & Black, 1993). Fortunately, heterogeneous flows that are dominated by the interaction of eddies with a mean shear are amenable to relatively simple closures, because the evolution of third-order cumulants, describing eddy-eddy interactions, can sometimes be neglected (Farrell & Ioannou, 2014). Statistical state dynamics, or direct statistical simulation (Tobias *et al.*, 2011; Ait-Chaalal *et al.*, 2016) has therefore been applied with success in simulations of planetary jets (Marston & Conover, 2008; Tobias & Marston, 2013) and wall-bounded shear-flow (Farrell *et al.*, 2016). Whilst strongly nonlinear systems, such as the model for Rayleigh-Bénard convection given by the Lorenz equations (Lorenz, 1963), require a more sophisticated treatment that accounts for the role of cumulants beyond second order, direct statistical simulation can nevertheless produce accurate predictions (Allawala & Marston, 2016).

The approach that we describe combines the desirable features of the statistical state equations with observations from direct simulation and classical adjoint techniques. In §2 we describe the problems associated with the adjoint analysis of chaotic systems, before deriving a well-conditioned adjoint operator from a system's cumulant equations in §3. In §4 we apply the approach to the sensitivity analysis of thermal convection via the Lorenz equations, and consider their extension to a 9-dimensional phase space in §5. Conclusions and suggestions for further work are made in §6.

2   The problem

Consider a dynamical system whose state, $\boldsymbol{Q}(t)$, evolves according to

$$\frac{\mathrm{d}\boldsymbol{Q}}{\mathrm{d}t} = \boldsymbol{F}(\boldsymbol{Q}, \boldsymbol{m}), \tag{1}$$

where $\boldsymbol{m}$ is a vector of system parameters. If the dynamical system (1) is chaotic then an understanding of the system's statistics becomes crucial. Fortunately, engineers and scientists are typically interested in a small subset of the possible statistics that can be obtained from (1). Unfortunately, they typically wish to understand how sensitively such statistics depend on *each* element of the parameter vector $\boldsymbol{m}$.

We focus our attention on the Lorenz equations as a specific example. Lorenz (1963) derived the following system of equations from a truncated description of Rayleigh-Bénard convection between hot and cold horizontal surfaces:

$$\frac{\mathrm{d}X}{\mathrm{d}t} = s(Y - X), \qquad \frac{\mathrm{d}Y}{\mathrm{d}t} = rX - Y - XZ, \qquad \frac{\mathrm{d}Z}{\mathrm{d}t} = XY - bZ. \tag{2}$$

The state $\boldsymbol{Q} = (X, Y, Z)$ describes the strength of the velocity field, the difference in temperature between ascending and descending fluid, and the strength of the horizontally averaged temperature with respect to a state of pure conduction, respectively. The parameters $\boldsymbol{m} = (s, b, r)$ are the Prandtl number, the aspect ratio of the resulting convection rolls and the Rayleigh number, respectively, the latter normalised with respect to a critical Rayleigh number.

We will focus on statistics $\boldsymbol{J}$ that correspond to a finite time average $\overline{\boldsymbol{L}}$ of a function
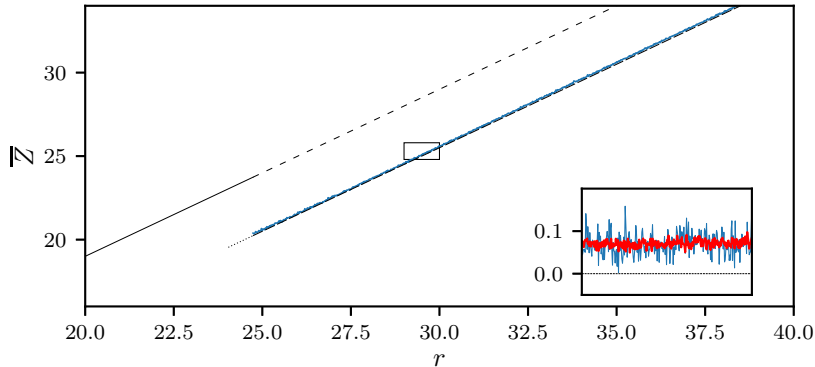
Figure 1: The dependence of the statistic $J = \overline{Z}$, as defined in (3), on the renormalised Rayleigh number r from a solution of the Lorenz equations (2) with $(s, b) = (10, 8/3)$. The results were obtained from integration along a statistically stationary trajectory for $\tau = 1000$ time units (blue/dark lines) and $\tau = 10000$ time units (red/light lines). The thin solid black line corresponds to the location $Z = r - 1$ of the stable fixed points of the system for $r < 24.7$ and to an upper bound of $\overline{Z}$ for $r \geq 24.7$. The dashed black straight line $\overline{Z} = r - 4.50$ is provided for reference. The dependence $\overline{Z} \propto r$ has been subtracted from the data that are displayed in the inset window.

$L(Q)$:

$$J[Q] \equiv \overline{L} = \frac{1}{\tau} \int\limits_0^\tau L(Q)\mathrm{d}t, \tag{3}$$

which depends implicitly on the parameters m via $Q(t)$. Under the assumption of ergodicity, the estimator J using the finite time average in (3) corresponds to a phase-average of the function L when $\tau \to \infty$. Following previous work on the sensitivity analysis of chaotic systems (Lea *et al.*, 2000; Wang, 2013), we will focus on $L \equiv Z$, such that $J \equiv \overline{Z}$ estimates the average amplitude of the horizontally-averaged temperature fluctuations. Figure 1 displays $\overline{Z}$ computed numerically from simulations of (2) with $s = 10$ and $b = 8/3$. In spite of the chaotic dynamics described by equation (2), $\overline{Z}$ appears to vary linearly with respect to $r > r_c$, for the values of r considered, where $r_c \approx 24.74$ is a critical value of r. At $r = r_c$ the two stable fixed points of the system, for the given values of s and b, become unstable. The value r_c marks the threshold of sustained chaotic behaviour, on which we focus, for almost all initial conditions, in contrast to the transient chaos that can be observed on an unstable chaotic set for $13.93 < r < 24.06$ (Yorke & Yorke, 1979). The oscillations in $\overline{Z}$ in figure 1 are due to the fact that $\overline{Z}$ is an estimator obtained from a finite time interval. Indeed, comparison of the statistic obtained over $\tau = 1000$ with that obtained over $\tau = 10000$ in figure 1, indicates that the oscillations reduce in amplitude as the length of the time interval increases.

2.1   *A finite difference approach*

An estimation of the Gâteaux derivative of a single functional J_j with respect to the ith component m_i of $\boldsymbol{m}$ is

$$\frac{\partial J_j}{\partial m_i} \approx \frac{\delta J_j}{\delta m_i} \equiv \frac{J_j[\boldsymbol{Q}(t|\boldsymbol{m}+\epsilon\boldsymbol{e}_i)] - J_j[\boldsymbol{Q}(t|\boldsymbol{m})]}{\epsilon}, \tag{4}$$

in which all elements of $\boldsymbol{e}_i$ are equal to zero, with the exception of the ith element, which is equal to 1. If ϵ is relatively large, then $\delta J_j/\delta m_i$ will not provide an accurate approximation to the local derivative. If, on the other hand, ϵ is relatively *small*, the non-smooth behaviour of J_j for finite-time averages evident in figure 1 suggests that we would need to obtain statistics over a correspondingly *large* time to obtain meaningful results (see e.g. Russo & Luchini, 2016). Moreover, the use of such an approach to obtain the sensitivity of J_j with respect to other system parameters requires the entire simulation to be run at least twice for each parameter m_i.

2.2   *The tangent linear equations*

Instead of by looking at finite differences between independent simulations, a functional's gradient can, in theory, be calculated exactly using information from a single simulation. We will outline a naive version of the method before explaining its problems in the case of chaotic systems. Assuming that it is well-defined, the derivative of a given functional J_j with respect to a given component m_i of $\boldsymbol{m}$ for a known trajectory $\boldsymbol{Q}(t)$ can be evaluated as

$$\frac{\partial J_j}{\partial m_i} = \frac{1}{\tau}\int\limits_0^\tau \frac{\partial L_j}{\partial \boldsymbol{Q}} \cdot \boldsymbol{q}_i \,\mathrm{d}t, \tag{5}$$

where $\boldsymbol{q}_i \equiv \mathrm{d}\boldsymbol{Q}/\mathrm{d}m_i$. By augmenting the time-dependent functions with boundary values:

$$\tilde{\boldsymbol{q}}_i \equiv \begin{bmatrix} \boldsymbol{q}_i(t) \\ \boldsymbol{q}_i(\tau) \end{bmatrix}, \qquad \tilde{\boldsymbol{g}}_j \equiv \begin{bmatrix} \dfrac{\partial L_j}{\partial \boldsymbol{Q}}(t) \\ \boldsymbol{0} \end{bmatrix}, \tag{6}$$

and defining an inner product $(\cdot,\cdot)$, based on (5), for the space to which the elements $\tilde{\boldsymbol{q}}_i$ and $\tilde{\boldsymbol{g}}_j$ belong, the functional's derivative can be expressed as (see e.g. Sewell, 1987)

$$\frac{\partial J_j}{\partial m_i} = \left(\tilde{\boldsymbol{g}}_j, \tilde{\boldsymbol{q}}_i\right). \tag{7}$$

In general, a penalty term that depends on the state trajectory's end point $\boldsymbol{Q}(\tau)$ can be added to (3), which would modify the $\boldsymbol{0}$ that appears in (6) for $\tilde{\boldsymbol{g}}_j$. The perturbation $\boldsymbol{q}_i$ satisfies the tangent linear equations, which are obtained by differentiating $\boldsymbol{F}$ with respect to m_i:

$$\mathcal{T}\tilde{\boldsymbol{q}}_i \equiv \begin{bmatrix} \dfrac{\mathrm{d}\boldsymbol{q}_i}{\mathrm{d}t} - \left(\dfrac{\partial \boldsymbol{F}}{\partial \boldsymbol{Q}}\right)\boldsymbol{q}_i \\ \boldsymbol{q}_i(0) \end{bmatrix} = \begin{bmatrix} \dfrac{\partial \boldsymbol{F}}{\partial m_i} \\ \boldsymbol{0} \end{bmatrix} \equiv \tilde{\boldsymbol{f}}_i, \tag{8}$$

where $\boldsymbol{q}_i(0)$ sets the perturbation of the initial condition to be zero. For the Lorenz equations the tangent linear system for $m_i = r$ is equivalent to

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} -s & s & 0 \\ r-Z & -1 & -X \\ Y & X & -b \end{pmatrix}\begin{pmatrix} x \\ y \\ z \end{pmatrix} + \begin{pmatrix} 0 \\ X \\ 0 \end{pmatrix}, \tag{9}$$

in addition to the initial condition $x(0) = y(0) = z(0) = 0$. Note that the X, Y and Z appearing in (9) are known, albeit time-dependent, variables. Once the perturbed trajectory $\boldsymbol{q}_i \equiv (x, y, z)$ is known, the derivative of *any* functional J_j with respect to m_i can in principle be calculated using (7) by changing $\boldsymbol{g}_j$. Calculation of the derivative of a given J with respect to a different parameter m is more difficult, because it requires us to find a different perturbed trajectory from (8) for use in (5). This motivates an alternative way of factorising (5), to obtain adjoint variables that describe the change in J with respect to a change in the constraints $\boldsymbol{F}$.

2.3   *The adjoint equations*

Introducing the adjoint variables $\boldsymbol{p}_j$ to enforce the the equations of motion $\boldsymbol{F}$, which act as constraints, results in

$$\frac{\partial J_j}{\partial m_i} = \left(\tilde{\boldsymbol{g}}_j, \tilde{\boldsymbol{q}}_i\right) + \langle \tilde{\boldsymbol{p}}_j, \tilde{\boldsymbol{f}}_i - T\tilde{\boldsymbol{q}}_i \rangle, \tag{10}$$

where $\langle \cdot, \cdot \rangle$ is an inner product for the dual space containing the elements

$$\tilde{\boldsymbol{p}}_j = \begin{bmatrix} \boldsymbol{p}_j(t) \\ \boldsymbol{p}_j(0) \end{bmatrix}, \qquad \tilde{\boldsymbol{f}}_i = \begin{bmatrix} \dfrac{\partial \boldsymbol{F}}{\partial m_i}(t) \\ \boldsymbol{0} \end{bmatrix}, \tag{11}$$

If $\tilde{\boldsymbol{p}}_j$ satisfies the adjoint equations:

$$T^\dagger \tilde{\boldsymbol{p}}_j \equiv \begin{bmatrix} -\dfrac{\mathrm{d}\boldsymbol{p}_j}{\mathrm{d}t} - \left(\dfrac{\partial \boldsymbol{F}}{\partial \boldsymbol{Q}}\right)^\dagger \boldsymbol{p}_j \\ \boldsymbol{p}_j(\tau) \end{bmatrix} = \begin{bmatrix} \dfrac{\partial L_j}{\partial \boldsymbol{Q}}(t) \\ \boldsymbol{0} \end{bmatrix} \equiv \tilde{\boldsymbol{g}}_j, \tag{12}$$

where $\dagger$ denotes the adjoint/transpose of an operator, then integration by parts of (10) results in

$$\frac{\partial J_j}{\partial m_i} = \langle \tilde{\boldsymbol{p}}_j, \tilde{\boldsymbol{f}}_i \rangle. \tag{13}$$

Note that the penalty term $\boldsymbol{0}$ in $\tilde{\boldsymbol{g}}_j$ corresponds to the initial condition $\boldsymbol{p}_j(\tau)$, and that the penalty term on $\boldsymbol{p}_j(0)$ in (13) corresponds to the initial condition $\boldsymbol{q}_i(0)$. Since $\boldsymbol{p}_j$ is equal to the derivative of J_j with respect to a change in the constraints, equation (13) enables us to readily compute the sensitivity of a given functional J_j with respect to *any* parameter m_i. To appreciate this, observe that (13) contains i, whereas (12) does not, and that the converse statement is true for (7) and (8). On the other hand, calculation of the the sensitivity of a different functional is difficult using the adjoint approach, because it would require a new solution of (12), which, like obtaining a solution to (8) is computationally demanding in comparison with the evaluation of (7) or (13).

Solutions to both the tangent system (8) and the adjoint system (12) grow without bound as the time τ in (3) increases. As pointed out by Thuburn (2005), the cause of the difference between the actual gradient and a gradient obtained from either the tangent or adjoint system is the fact that the operation of time averaging over $\tau \to \infty$ does not, in general, commute with the finite difference of an infinite time average over an interval of $\epsilon \to 0$:

$$\frac{\partial J}{\partial m_i} = \lim_{\epsilon \to 0} \lim_{\tau \to \infty} \frac{\delta J}{\delta m_i} \neq \lim_{\tau \to \infty} \lim_{\epsilon \to 0} \frac{\delta J}{\delta m_i}. \tag{14}$$

The finite difference of the functional J_j does not converge uniformly to the sought-after derivative for all integration times τ and, therefore, neither do (7) nor (13).

An approximation to $\partial_{m_i} J_j$ can be obtained if (12) is integrated over relatively short time intervals (Lea *et al.*, 2000). On the other hand, if finite differences are employed using equation (4), then the minimal time τ over which accurate statistics can be obtained is determined by $\epsilon \ll 1$. An accurate finite difference approximation requires a small value of ϵ, which requires a large value of τ (see e.g. Russo & Luchini, 2016). Thus, approximate gradients can be obtained by using the tangent linear equations or finite differences, provided that small or large time intervals are used, respectively.

That it is crucial to take the limit $\tau \to \infty$ *before* analysing derivatives suggests that sensitivity analysis of the equations governing the statistics of the process might result in a more tractable problem. At the expense of introducing additional unknowns, we therefore focus on obtaining adjoint information for the equations satisfied by the system's cumulants.

3   The cumulant equations and their closure

The equations that govern the behaviour of cumulants provide a means of establishing the leading-order relationships between the statistics of a chaotic attractor. These relationships constrain the response of statistics to changes in parameters. The cumulants and their dynamics have a natural hierarchy and can be derived in a systematic way from the equations that govern individual trajectories.

3.1   *The cumulant generating functional*

The cumulants $\boldsymbol{U}$ of a dynamical system can be defined in terms of a cumulant generating functional $\log \overline{\psi}$:

$$U_{\alpha_1 \alpha_2 \ldots \alpha_d} = (-\imath)^{|\alpha|} \left. \frac{\partial^\alpha}{\partial \boldsymbol{P}^\alpha} \log \overline{\psi} \right|_{\boldsymbol{P}=0}, \tag{15}$$

where $\alpha = (\alpha_1, \alpha_2 \ldots, \alpha_d)$ is a multi-index for the system of d time-dependent variables $\boldsymbol{Q}(t)$, and

$$\psi \equiv \exp\left(\imath \, P_i Q_i(t)\right), \tag{16}$$

where ψ is the Hopf generating functional (see e.g. Hopf, 1952; Frisch, 1995) and $\imath = \sqrt{-1}$. The over-bar in (15) denotes the finite time average defined in (3), which we assume converges to a phase average when the duration the averaging interval $\tau \to \infty$. Due

to the logarithm in (15), a cumulant, unlike a moment, derived from the sum of two independent random variables, is equal to the sum of their respective cumulants. This commutativity is related to the fact that cumulants isolate the interdependence of random variables without including the effects of correlations between statistics of lower order. For example, according to (15) and (16),

$$\overline{Z^3} = U_{003} + 3U_{002}U_{001} + U_{001}^3, \tag{17}$$

where the coefficients of the three terms on the right-hand side correspond to the number of ways of partitioning a multiset of three (identical) elements into (a) a single multiset of three; (b) a multiset of two and a set of one; (c) three sets of one. In this respect, cumulants are the atoms of which moments are comprised, and therefore have simpler algebraic properties than the latter. Further examples of the decomposition of moments into cumulants include $\overline{XZ} = U_{101} + U_{100}U_{001}$, and

$$\overline{XYZ^2} = \overbrace{U_{112}}^{\text{4th order}} + \overbrace{2U_{111}U_{001} + U_{100}U_{012} + U_{010}U_{102}}^{\text{3rd and 1st order}} + \overbrace{U_{110}U_{002} + 2U_{101}U_{011}}^{\text{2nd order}}$$
$$+ \underbrace{U_{100}U_{010}U_{002} + 2U_{100}U_{001}U_{011} + 2U_{010}U_{001}U_{101}}_{\text{2nd and 1st order}} + \underbrace{U_{100}U_{010}U_{001}^2}_{\text{1st order}}, \tag{18}$$

in which the grouped terms correspond to a summation over the different ways that the multiset of four elements $\{X, Y, Z, Z\}$ can be partitioned into subsets of a given cardinality. The connection between cumulants and moments is discussed in more detail in appendix A, in which it is helpful to compare (18) with (47).

Noting from (16) that $\boldsymbol{Q}$ plays the role of $-\imath\partial_{\boldsymbol{P}}$, the Hopf function ψ satisfies the *linear* equation

$$\imath\frac{\partial\psi}{\partial t} = -P_i F_i\left(-\imath\frac{\partial}{\partial\boldsymbol{P}}, \boldsymbol{m}\right)\psi. \tag{19}$$

The original d nonlinear equations $\boldsymbol{F}$ from (1) are recovered by differentiating (19) with respect to the vector $\boldsymbol{P}$. Associated with the original system (1) are an infinite hierarchy of cumulant equations,

$$\boldsymbol{H}(\boldsymbol{U}, \boldsymbol{m}) = 0, \tag{20}$$

which are obtained, under the assumption of ergodicity, by averaging (19) to obtain a stationary equation for $\overline{\psi}$, in which the equation for a given cumulant U_α corresponds to the coefficient of $\boldsymbol{P}^\alpha$. Readers are referred to Frisch (1995) for further details.

Despite the fact that they do not form a closed system, the cumulant equations provide useful information. For example, as noted by Knobloch (1979), the cumulant equations for the Lorenz system indicate that

$$\overline{Z} = \frac{\overline{X^2}}{b} = r - 1 - \frac{1}{s^2\overline{Z}}\frac{\mathrm{d}\overline{X^2}}{\mathrm{d}t} - \frac{\overline{Z^2} - \overline{Z}^2}{\overline{Z}}, \tag{21}$$

which, since $\overline{Z^2} \geq \overline{Z}^2$, implies that $0 \leq \overline{Z} \leq r - 1$ in a statistically steady state.

If the original system evolves on a d-dimensional phase space then, ignoring symmetries in the governing equations, the number of cumulants at order $j = |\alpha|$ is equal to the number
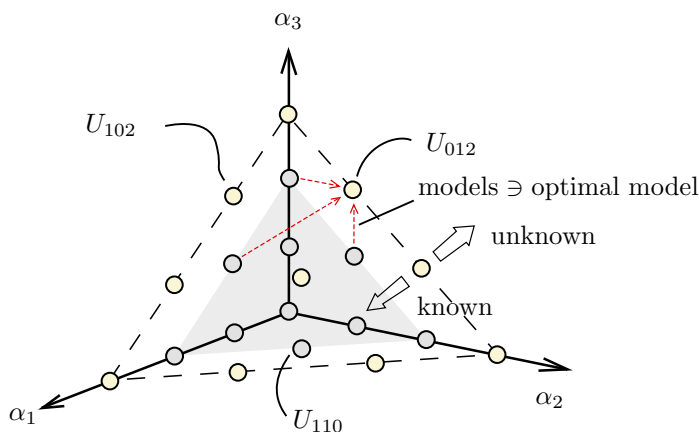
Figure 2: Cumulant space for $d = 3$ indexed with $\alpha = (\alpha_1, \alpha_2, \alpha_3)$ and truncation at order $N = 2$. For systems, such as (2), with quadratic nonlinearities, a model is the specification of the dependence of cumulants of order $N + 1$ on those of lower order.

of ways that j indistinguishable objects can be assigned to d sets; hence the number of cumulants up to and including those of order N is

$$\sum_{j=1}^{N} \binom{j + d - 1}{d - 1},\tag{22}$$

as illustrated in figure $2(a)$. Known symmetries of a system reduce the number of independent unknown cumulants. In the case of the Lorenz equations (2), for which $\boldsymbol{F}$ is invariant under the mapping $(X, Y) \mapsto (-X, -Y)$, a cumulant $U_{\alpha_1 \alpha_2 \alpha_3}$ for which $\alpha_1 + \alpha_2$ is odd, is necessarily equal to zero. In the case of $d = 3$, of the 34 and 55 available cumulants up to order $N = 4$ and $N = 5$, 18 and 27, respectively, are non zero. Closures of the cumulant hierarchy aim to strike a balance between the incorporation of additional physics from nonlinear interactions and keeping the number of unknowns small.

3.2   *Adjoint cumulant dynamics*

As described in §2, if one wishes to differentiate a vector-valued function with respect to vector-valued input one can employ one of two dual approaches. Each approach uses the chain rule: one forwards and the other backwards. The first approach is to propagate derivatives with respect to a single input parameter 'up' the computational graph, to find the derivatives of all output variables. The second approach is to propagate derivatives of a single output variable 'down' the computational graph with respect to all input variables. The efficiency of the two methods depends on the number of input parameters compared with the number of output variables.

Regarding $\boldsymbol{J}[\boldsymbol{U}(\boldsymbol{m})]$, originally defined in (3), as a functional of an infinite hierarchy of cumulants,

$$\frac{\partial \boldsymbol{J}}{\partial \boldsymbol{m}} = -\underbrace{\frac{\partial \boldsymbol{J}}{\partial \boldsymbol{U}} \overbrace{\left(\frac{\partial \boldsymbol{H}}{\partial \boldsymbol{U}}\right)^{-1} \frac{\partial \boldsymbol{H}}{\partial \boldsymbol{m}}}^{-\boldsymbol{u}}}_{\boldsymbol{v}},\tag{23}$$

where the cumulant perturbation $\boldsymbol{u}$ and the corresponding adjoint variables $\boldsymbol{v}$ are defined according to

$$\boldsymbol{u} \equiv \frac{\partial \boldsymbol{U}}{\partial \boldsymbol{m}} = -\left(\frac{\partial \boldsymbol{H}}{\partial \boldsymbol{U}}\right)^{-1}\frac{\partial \boldsymbol{H}}{\partial \boldsymbol{m}}, \qquad \boldsymbol{v} \equiv \frac{\partial \boldsymbol{J}}{\partial \boldsymbol{H}} = \frac{\partial \boldsymbol{J}}{\partial \boldsymbol{U}}\left(\frac{\partial \boldsymbol{H}}{\partial \boldsymbol{U}}\right)^{-1}. \tag{24}$$

As discussed in §2, if the problem involves more functionals than parameters, it is computationally preferable to solve the tangent linear system to find $\boldsymbol{u}$ before evaluating (23). If, on the other hand, the problem contains more parameters than functionals, then it is computationally preferable to find the adjoint variables $\boldsymbol{v}$ before evaluating (23). For a given functional J_j and a given parameter m_i, the two alternatives can be expressed as

$$\frac{\partial J_j}{\partial m_i} = \begin{cases} (\boldsymbol{g}_j, \boldsymbol{u}_i) & \text{s.t.} \quad \mathcal{T}\boldsymbol{u}_i = \boldsymbol{f}_i, \\ \langle \boldsymbol{v}_j, \boldsymbol{f}_i \rangle & \text{s.t.} \quad \mathcal{T}^\dagger \boldsymbol{v}_j = \boldsymbol{g}_j, \end{cases} \tag{25}$$

where

$$\mathcal{T} = \frac{\partial \boldsymbol{H}}{\partial \boldsymbol{U}}, \qquad \boldsymbol{f}_i = -\frac{\partial \boldsymbol{H}}{\partial m_i}, \qquad \boldsymbol{g}_j = \frac{\partial J_j}{\partial \boldsymbol{U}}, \tag{26}$$

We focus on the adjoint problem of determining the sensitivity of a single functional J (we omit the subscript j hereafter) with respect to a potentially large number of unknown parameters. Unlike the systems (8) and (12), whose validity relies on the commutation of time averaging and differentiation with respect to m_i, (23) works with time averaged variables directly.

3.3   *Building the cumulant operator*

If the original system of equations (1) contains nonlinear terms then the equations for the cumulants of order j will depend on cumulants of order $j+1$ and higher, depending on the degree of nonlinearity. For the Lorenz equations (2), and indeed the quadratic equations governing fluid mechanics more generally, cumulants of order j do not have a dependence on cumulants whose order is higher than $j + 1$. It is nevertheless necessary to close the problem, as illustrated by the shape of the tangent linear operator:

$$\begin{bmatrix} \mathcal{T}^{(11)} & \mathcal{T}^{(12)} & 0 & 0 & \cdots \\ \mathcal{T}^{(21)} & \mathcal{T}^{(22)} & \mathcal{T}^{(23)} & 0 & \cdots \\ \vdots & \vdots & \vdots & & \end{bmatrix} \begin{pmatrix} \boldsymbol{u}_i^{(1)} \\ \boldsymbol{u}_i^{(2)} \\ \boldsymbol{u}_i^{(3)} \\ \vdots \end{pmatrix} = \boldsymbol{f}. \tag{27}$$

Here $\boldsymbol{u}_i^{(j)}$ represents perturbations $\partial_{m_i} U^{(j)}$ to the cumulants of order $j = |\alpha|$. According to (22), each operator $\mathcal{T}^{(ij)}$ has the shape

$$\text{shape}(\mathcal{T}^{(ij)}) = \binom{i+d-1}{d-1} \times \binom{j+d-1}{d-1}. \tag{28}$$

For the Lorenz equations,

Figure 3: (a) Modification of the functional due to the inner product of unclosed cumulants $\boldsymbol{u}^{(N+1)}$ and the weighting factor $\boldsymbol{g}^{(N+1)}$. (b) Local observation of the underlying functional $J[\boldsymbol{U}(\boldsymbol{m})]$ (red circle) and gradient ∇J_* from the optimal model as an approximation of the underlying exact gradient ∇J.

$$[\mathcal{T}^{(11)}, \mathcal{T}^{(12)}] = \begin{bmatrix} s & -s & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ r - \overline{Z} & -1 & -\overline{X} & 0 & 0 & -1 & 0 & 0 & 0 \\ \overline{Y} & \overline{X} & b & 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}. \tag{29}$$

Whilst the tangent linear system is *under determined*, the adjoint system $\mathcal{T}^\dagger \boldsymbol{v} = \boldsymbol{f}$ is *over determined*. The overall properties of the system can be be seen in the self-adjoint problem that combines the tangent linear and adjoint operators. With the equations for perturbations to the first order cumulants (e.g. $\overline{X}, \overline{Y}$ and $\overline{Z}$), one finds

$$\begin{bmatrix} 0 & \mathcal{T}^{(11)} & \mathcal{T}^{(12)} \\ \mathcal{T}^{(11)\dagger} & 0 & 0 \\ \mathcal{T}^{(12)\dagger} & 0 & 0 \end{bmatrix} \begin{pmatrix} \boldsymbol{v}^{(1)} \\ \boldsymbol{u}_i^{(1)} \\ \boldsymbol{u}_i^{(2)} \end{pmatrix} = \begin{pmatrix} \boldsymbol{f}_i^{(1)} \\ \boldsymbol{g}^{(1)} \\ \boldsymbol{g}^{(2)} \end{pmatrix}. \tag{30}$$

If the functional J, and therefore the vector $\partial_{\boldsymbol{U}^{(1)}} J \equiv \boldsymbol{g}^{(1)}$, is specified then one can solve for the adjoint variables $\boldsymbol{v}^{(1)}$ according to the second row of (30). However, a consistency requirement for the extended system (30) to possess a solution is that $\boldsymbol{g}^{(2)} = \mathcal{T}^{(12)\dagger} \boldsymbol{v}^{(1)} \neq 0$, in general. We are therefore not at liberty to choose the functional J arbitrarily, because it will automatically contain a contribution scaled by $\boldsymbol{g}^{(2)}$ from the unclosed perturbations $\boldsymbol{u}_i^2$.

The vacuous consequence of using (30) is that only functionals whose value can be determined identically from the original cumulant equations, such as equation (21), can be determined exactly. For (30) to yield novel information an assumption is required about the response of the unknown cumulant perturbations $\boldsymbol{u}_i^{(2)}$. The simplest, albeit naive, approach is to assume that $(\boldsymbol{g}^{(2)}, \boldsymbol{u}_i^{(2)}) = 0$, which corresponds to the unknown high-order perturbations $\boldsymbol{u}_i^{(2)}$ being either zero or orthogonal to the weighting vector $\boldsymbol{g}^{(2)}$, as illustrated in figure 3(a). More generally, taking the system (30) as an example, a closure corresponds to the specification of $(\boldsymbol{g}^{(2)}, \boldsymbol{u}_i^{(2)})$ in terms of the the retained cumulant

sensitivities $\boldsymbol{u}_i^{(1)}$. If $\boldsymbol{m}$ belongs to a three-dimensional parameter space, then truncation at order N entails three assumptions, determining $(\boldsymbol{g}^{(N+1)}, \boldsymbol{u}_i^{(N+1)})$ for $i = 1, 2, 3$. In this respect, the number of required assumptions is independent of the order N at which a closure is invoked, which arguably makes finding a suitable closure for sensitivity analysis less onerous than finding a suitable closure for the original cumulant equations.

3.4 Closure

As illustrated in figure 2, to obtain a closed system of cumulant equations one needs to make an assumption about how the highest-order cumulants are related to those of lower order and, therefore, the way in which they depend on the problem's parameters. One approach is to assume that cumulants whose order is higher than N are not affected, or respond sufficiently slowly, to changes in the problem's parameters, which is a sufficient condition for $(\boldsymbol{g}^{(N+1)}, \boldsymbol{u}^{(N+1)}) = 0$. For $N = 2$ in a system with quadratic nonlinearities, this approach is consistent with the assumption that the probability distribution of the underlying process is Gaussian and is therefore completely determined by its cumulants of first and second order (Frisch, 1995).

As outlined in section §3.3, for sensitivity analysis the implications of discarding cumulants beyond a certain order are weaker than those associated with direct simulation of the truncated equations. For example, truncation of the cumulant equations at order $N = 3$ and assuming that $\boldsymbol{U}^{(4)} \equiv 0$, produces non-realisable statistics (Kraichnan, 1980), leading to a negative energy spectrum in turbulence (Ogura & Phillips, 1962). From the perspective of sensitivity analysis, however, the orthogonality condition $(\boldsymbol{g}^{(4)}, \boldsymbol{u}^{(4)}) = 0$ does not necessarily imply that $\boldsymbol{U}^{(4)} \equiv 0$. Similarly, $(\boldsymbol{u}^{(3)}, \boldsymbol{g}^{(3)}) = 0$ does not necessarily imply that the process is Gaussian. It is nevertheless important to note that the behaviour of higher-order cumulants in a Gaussian distribution is a special case, because probability distributions possessing non-zero cumulants at order $N_* > 2$, followed by zero cumulants at all orders $N > N_*$, do not exist (Lukacs, 1970, p. 223).

One can discard cumulants of order higher than N and model their effects with a forcing function such as $\boldsymbol{M}^{(N)}$, which, in general, will depend on a vector $\boldsymbol{\mu}$ of unknown parameters:

$$\boldsymbol{H}^{(N)}(\pi_N \boldsymbol{U}, \boldsymbol{m}) = \boldsymbol{R}^{(N)} + \boldsymbol{M}^{(N)}(\pi_N \boldsymbol{U}, \boldsymbol{\mu}), \tag{31}$$

where $\boldsymbol{R}^{(N)}$ represents the residuals arising from the truncation and $\pi_N \boldsymbol{U}$ is the projection that sets the value of cumulants whose order exceeds N to zero. Assuming that the residual $\boldsymbol{R}^N$ can be made small with a suitable choice of $\boldsymbol{M}^{(N)}$, and that for a given $\boldsymbol{M}^{(N)}$, $\boldsymbol{R}^{(N)}$ does not depend on $\boldsymbol{m}$, the tangent linear equations at order N are

$$\left(\frac{\partial \boldsymbol{H}^{(N)}}{\partial \boldsymbol{U}} - \frac{\partial \boldsymbol{M}^{(N)}}{\partial \boldsymbol{U}} \right) \boldsymbol{u}^{(N)} = -\frac{\partial \boldsymbol{H}^{(N)}}{\partial \boldsymbol{m}}. \tag{32}$$

A key assumption underlying the use of (32) as a model for the tangent linear behaviour of the system is that the model parameters $\boldsymbol{\mu}$ in (31) exhibit a weak dependence on the problem parameters $\boldsymbol{m}$ (hence $\partial_{\boldsymbol{m}} \boldsymbol{M}$ is not included in (32)), which is consistent with the assumption that $\boldsymbol{R}^{(N)} = 0$ in the vicinity of $\boldsymbol{m}$. Utilising (32) for truncation at $N = 3$ in the sensitivity analysis of a system with quadratic nonlinearities, under the assumption that $\boldsymbol{M}^{(N)}$ depends only on the highest retained cumulants $\boldsymbol{U}^{(N)}$, yields

$$
\begin{bmatrix}
0 & 0 & 0 & T^{(11)} & T^{(12)} & 0 \\
0 & 0 & 0 & T^{(21)} & T^{(22)} & T^{(23)} \\
0 & 0 & 0 & T^{(31)} & T^{(32)} & T^{(33)} - M^{(33)} \\
T^{(11)\dagger} & T^{(21)\dagger} & T^{(31)\dagger} & 0 & 0 & 0 \\
T^{(12)\dagger} & T^{(22)\dagger} & T^{(32)\dagger} & 0 & 0 & 0 \\
0 & T^{(23)\dagger} & T^{(33)\dagger} - M^{(33)\dagger} & 0 & 0 & 0
\end{bmatrix}
\begin{pmatrix}
\boldsymbol{v}^{(1)} \\ \boldsymbol{v}^{(2)} \\ \boldsymbol{v}^{(3)} \\ \boldsymbol{u}^{(1)} \\ \boldsymbol{u}^{(2)} \\ \boldsymbol{u}^{(3)}
\end{pmatrix}
=
\begin{pmatrix}
\boldsymbol{f}^{(1)} \\ \boldsymbol{f}^{(2)} \\ \boldsymbol{f}^{(3)} \\ \boldsymbol{g}^{(1)} \\ \boldsymbol{g}^{(2)} \\ \boldsymbol{g}^{(3)}
\end{pmatrix},
\tag{33}
$$

where

$$
M^{(ij)} \equiv \frac{\partial \boldsymbol{M}^{(i)}}{\partial \boldsymbol{U}^{(j)}}.
\tag{34}
$$

The closed system of extended equations (33) is, in general, invertible and therefore provides a set of solutions for the adjoint variables $\boldsymbol{v}$ for a specified set of weights $\boldsymbol{g}$. Without selecting the model parameters $\boldsymbol{\mu}$, inversion of the adjoint operator $(T - M)^{\dagger}$ yields a fan of gradients, as indicated in figure 3(b). The determination of a unique gradient from the fan requires the selection of an optimal set of model parameters $\boldsymbol{\mu}_*$. For example, the optimal parameters could be chosen to minimise $\boldsymbol{R}^{(N)}$ according to statistical observations from a direct simulation:

$$
\boldsymbol{\mu}_* = \arg\min_{\boldsymbol{\mu}} \left\| \boldsymbol{H}^{(N)} - \boldsymbol{M}^{(N)} \right\|.
\tag{35}
$$

Using local observational data the procedure of obtaining sensitivity information can therefore be freed from tunable parameters once a suitable class of models has been selected.

The extent to which it is necessary to include cumulants of order greater than N for sensitivity calculations depends on the role they play in maintaining the statistical equilibrium defined by (20). Although the truncation of the cumulants at second order yields realisable statistics, the second order cumulants alone will in general not be capable of describing the fully nonlinear features of a flow (Frisch, 1995). As described above, inclusion of the third-order cumulants (the quasi-normal approximation) without accounting for the fourth-order cumulants is problematic in simulations, because the latter play a crucial role in damping the third-order cumulants (Bohr *et al.*, 2005). Therefore, a popular choice, known as the Eddy Damped Quasi-Normal Markovian approximation (see e.g. Leith & Kraichnan, 1972), is to truncate the cumulants at third order and to include a damping term to account for the discarded fourth-order cumulants:

$$
\boldsymbol{M}^{(3)} = \mu \boldsymbol{U}^{(3)}; \quad \text{hence } M = \left[0, 0, \mu I \right].
\tag{36}
$$

When $\mu \to \infty$ the cumulants of order $N = 3$ become increasingly damped and the closure corresponds to a truncation at $N = 2$; when $\mu \to 0$ the closure corresponds to truncation at $N = 3$ (Allawala & Marston, 2016). The eddy-damping parameter therefore produces a fan of possible functional gradients, as illustrated in figure 3(b).

The optimal value of μ that minimises the size of the normed residual $\|\boldsymbol{R}^{(3)}\|$ is

$$
\mu_* = \frac{(\boldsymbol{U}^{(3)}, \boldsymbol{H}^{(3)})}{\|\boldsymbol{U}^{(3)}\|^2},
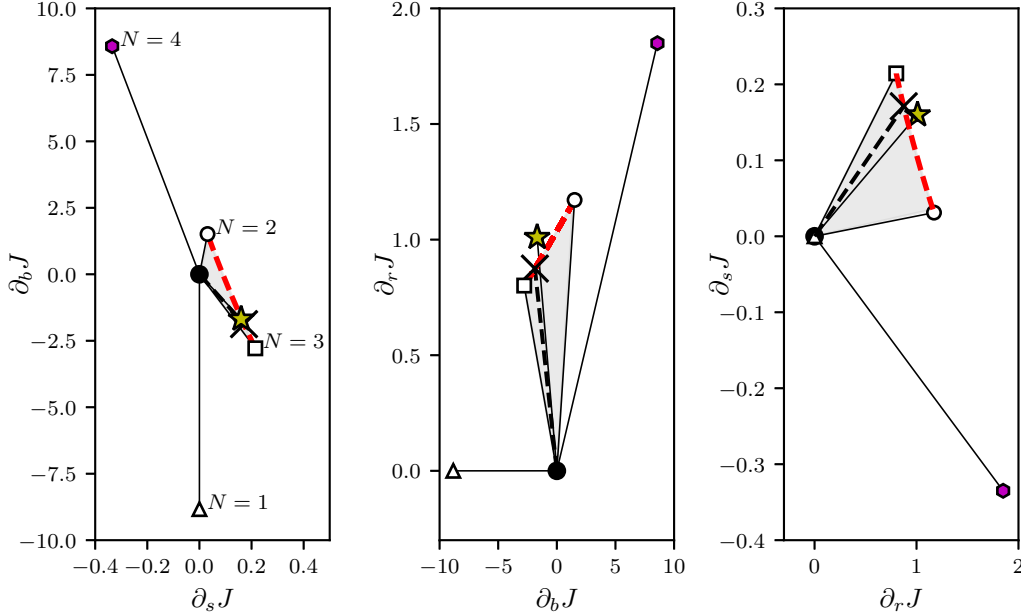\tag{37}
$$

Figure 4: Direction of the gradient $\nabla J = \nabla \overline{Z}$ evaluated at $m = (10, 8/3, 28)$ using the hierarchy of cumulant equations truncated at order $N = 1, 2, 3, 4$. The star corresponds to the observed gradient reported in Wang (2013) and $\times$ corresponds to the optimised eddy-damped model. Statistics were obtained over $\tau = 10^6$ dimensionless time units. Points along the thick dashed red line that connects $N = 2$ with $N = 3$ were obtained by determining $\nabla \overline{Z}$ for different values of $\mu \in (0, \infty)$.

which enables the optimal functional gradient to be determined according to

$$
\frac{\partial J}{\partial m_i} = \left\langle \left(\begin{bmatrix} T^{(11)\dagger} & T^{(21)\dagger} & T^{(31)\dagger} \\ T^{(12)\dagger} & T^{(22)\dagger} & T^{(32)\dagger} \\ 0 & T^{(23)\dagger} & T^{(33)\dagger} \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -\mu_* I \end{bmatrix} \right)^{-1} \boldsymbol{g}, \boldsymbol{f}_i \right\rangle. \tag{38}
$$

The procedure described in this section consists of identifying the order N at which the cumulant hierarchy should be truncated, before selecting a subclass of possible models for the unknown cumulants. The optimal parameters $\boldsymbol{\mu}_*$, and therefore optimal gradient ∇J_* in figure 3(b), can be determined by minimising the residual between statistics from direct simulation and the corresponding model prediction according to equation (35).

4   Two-dimensional convection (the Lorenz equations, $d = 3$)

4.1   *Truncation of the cumulant hierarchy*

To test the method for obtaining functional gradients described in §3.4, we collect statistics from direct simulations of the Lorenz equations (2). We focus on the statistically stationary state produced by parameters $(s, b, r) = (10, 8/3, 28)$, which is well documented and was the state chosen for the sensitivity analysis of Wang (2013). The dynamical equations are integrated using the DOPRI5 explicit Runge-Kutta method in Python's SciPy library. To

check convergence of the computed cumulants the time τ used to define the time average (3) was varied from $\tau = 10^3$ to $\tau = 10^6$. To allow for transient behaviour, the initial time used in the simulations is -100. Integrals such as (3) were computed using a trapezium rule over the discrete points obtained from the simulations.

Gradients of the functional $J = \overline{Z}$ are displayed in figure 4, which shows the projection of the gradient vector ∇J onto two-dimensional planes. The symbols denote the gradients that are obtained by truncating the cumulant hierarchy at order $N = 1, 2, 3, 4$, without modelling the discarded cumulants. Truncation of the tangent linear system at $N = 1$ yields an inaccurate representation of the gradient of $\overline{Z}$. The response of the second order cumulants to changes in the parameters is evidently significant and therefore the assumption that their dependence on parameters is identically zero (or, more generally, orthogonal to $\boldsymbol{g}^{(2)}$, as described in §3.3) produces poor predictions. Truncation of the tangent linear system at $N = 2$ also yields a poor approximation of $\nabla\overline{Z}$, particularly $\partial_b\overline{Z}$, but one that is an improvement in comparison with truncation at $N = 1$. As discussed in §3.4, in shearless turbulence the effect on eddies of eddy-eddy interactions, captured by the third order cumulants (Farrell & Ioannou, 2014), is expected to play a crucial role in maintaining statistical equilibrium in the case of the Lorenz equations. Indeed, the third order cumulants play a dynamically important role in determining the response of the Lorenz system to parametric changes, and figure 4 shows that their retention yields a reasonable approximation of $\nabla\overline{Z}$.

Truncation of the cumulant equations at $N = 4$ yields a poor approximation to $\nabla\overline{Z}$, which illustrates the need to find a compromise between the efficiency and simplicity of truncation at relatively low order and the additional physics that is captured by higher-order cumulants. In the absence of physical justification, truncation at higher order, rather than lower order, does not necessarily imply an improved estimation of the behaviour of the retained cumulants. Indeed, as noted in §3.4, distributions with cumulants that are non-zero up to order N, followed by cumulants that are zero above order N, are not realisable for $N > 2$. In this respect, it is perhaps not surprising that the fourth order approximation shown in figure 4 is inaccurate.

4.2   *Error analysis*

The difference between the approximation $\sum_{j=1}^{N}\langle\boldsymbol{v}^{(j)}, \boldsymbol{f}_i^{(j)}\rangle$ and the observed gradient $\nabla\overline{Z}$ that was obtained by truncating the cumulant hierarchy (depicted in figure 4 with a star) can be understood by inspecting the derivatives of the discarded cumulants. As discussed at the end of §3.3, the error associated with the i^{th} component of the gradient $\nabla\overline{Z}$ for truncation at order N is $(\boldsymbol{g}^{(N+1)}, \boldsymbol{u}_i^{(N+1)})$, where $\boldsymbol{u}_i^{(N+1)}$ are the perturbations of the neglected cumulants, and $\boldsymbol{g}^{(N+1)}$ determines the influence they have on the functional in question:

$$\frac{\partial J}{\partial m_i} = \underbrace{\sum_{j=1}^{N}\langle\boldsymbol{v}^{(j)}, \boldsymbol{f}_i^{(j)}\rangle}_{\text{approximation}} - \underbrace{\left(\boldsymbol{g}^{(N+1)}, \boldsymbol{u}_i^{(N+1)}\right)}_{\text{error}}. \tag{39}$$

We focus on the error associated with the derivative of $\overline{Z}$ with respect to r (i.e. $i = 3$), and display $\boldsymbol{g}^{(N+1)}$ and $\boldsymbol{u}_3^{(N+1)}$ for $N = 1, 2, 3$ and 4 in figures 5 and 6. We restrict atten-
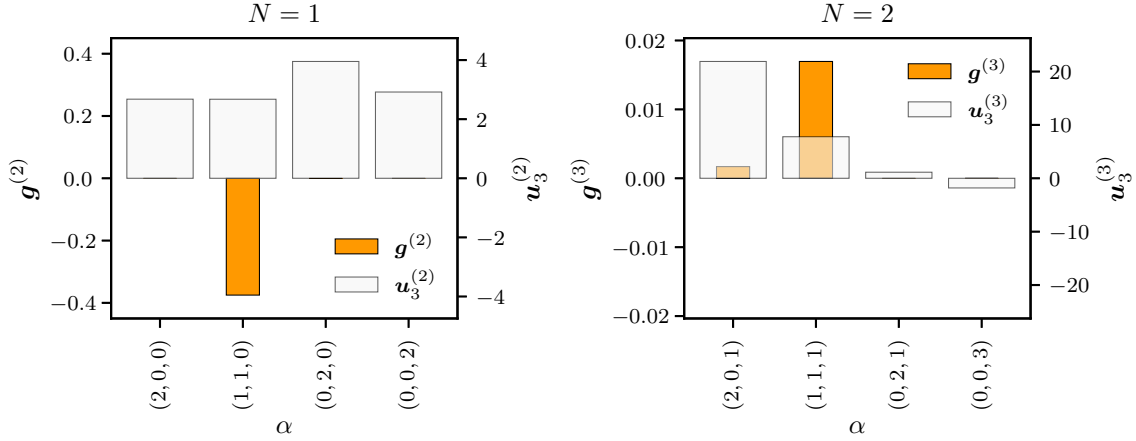
Figure 5: The constituent parts of the error $(\boldsymbol{g}^{(N+1)}, \boldsymbol{u}_3^{(N+1)})$ for truncation at order $N = 1$ (left) and $N = 2$ (right). The values used to create this figure can be found in tables 3 and 4 in appendix B.

tion to non-zero cumulants using the symmetry arguments made in §3.1. The gradients were determined by analysing statistics from 256 simulations employing values of r distributed uniformly over a unit interval centred on $r = 28$. Further details are provided in appendix B.

The error associated with truncation at $N = 1$ in figure 5 is entirely due to the behaviour of the cumulant $\overline{XY} = U_{110}$. The remaining cumulants, for which the corresponding values of $\boldsymbol{g}^{(2)}$ are zero, do not contribute to the error, as can be seen directly from the second row of the cumulant equations in equation (29). At order $N = 2$ in figure 5, the dominant contribution to the error comes from the response of U_{111}, and at order $N = 3$, shown in figure 6, it comes predominantly from U_{310} and U_{130}, which are related to the moments $\overline{XY^3}$ and $\overline{X^3Y}$. For truncation at order $N = 4$, the perturbations in the discarded cumulants are large $O(10^3)$, with figure 6 indicating that the dominant contribution to the error comes from U_{311}, which is related to the moment $\overline{X^3YZ}$. The effect on the error of the growing sensitivity and number of discarded cumulants is, to a limited extent, compensated by their diminishing influence on the gradient $\partial_r \overline{Z}$, as evidenced by the relatively small values of $\boldsymbol{g}^{(5)}$ in figure 6.

A summary of the truncation errors obtained at each order is provided in table 1. Obtaining accurate observations of the sensitivity of fifth-order statistics from the Lorenz attractor is challenging, because it requires the use of relatively large intervals for time averaging. The approximate equality between the third and fourth columns of table 1 nevertheless indicates that the sum of the inferred gradient $\sum_{j=1}^{N} \langle \boldsymbol{v}^{(j)}, \boldsymbol{f}_i^{(N)} \rangle$ and the error $-(\boldsymbol{g}^{(N+1)}, \boldsymbol{u}_3^{(N+1)})$ agrees with $\partial_r \overline{Z}$, and therefore satisfies the original cumulant equations to within 2%. At orders $1, 2$ and 3 the difference between the third and fourth columns of table 1 implies that the cumulant equations are satisfied to within approximately 0.1%.
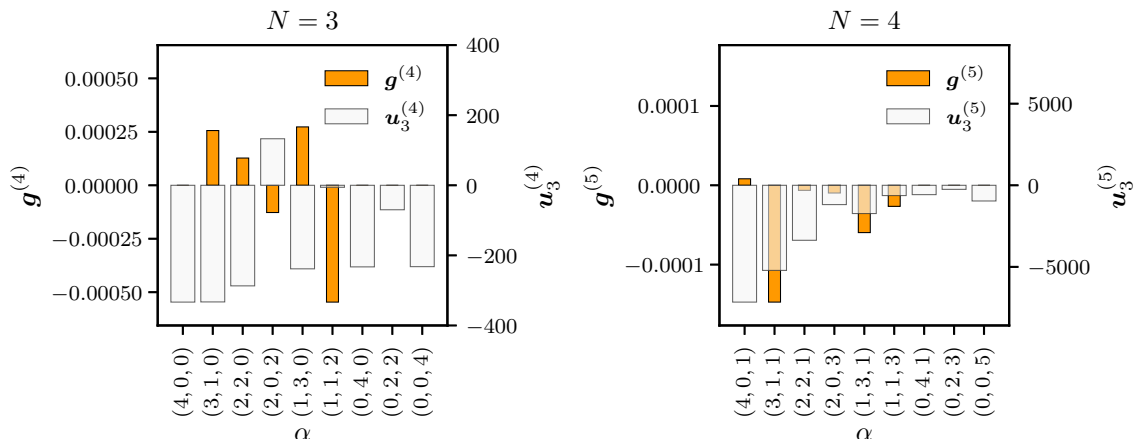
Figure 6: The constituent parts of the error $(\boldsymbol{g}^{(N+1)}, \boldsymbol{u}_3^{(N+1)})$ for truncation at order $N = 3$ (left) and $N = 4$ (right). The values used to create this figure can be found in tables 5 and 6 in appendix B.

4.3   *Optimal closure*

In addition to the relatively simple truncations discussed in §§4.1-4.2, corresponding to assumption that $(\boldsymbol{g}^{(N+1)}, \boldsymbol{u}^{(N+1)}) = 0$, the projections in figure 4 also display the gradients that are obtained by varying the eddy-damping parameter μ described in §3.4. The resulting family of gradients produce a fan of gradient vectors lying between the limit points associated with second-order truncation ($\mu \to \infty$) and the third-order truncation ($\mu \to 0$). A single member of the family corresponds to the eddy damping that is optimal, in the sense of equation (35), with respect to observations. Although the optimal eddy damping μ_* yields a gradient that is close to the observed gradient, figure 4 indicates that other values of μ would yield a slightly improved prediction. The reason for this is that the parameter that minimises the residual of the difference between the cumulant equations and the observations is not necessarily that which minimises the difference between the predicted and observed gradients of a given functional, and therefore typifies the difficulties of deriving gradients from a single set of statistics.

Figure 7 displays orthogonal slices through the functional $\overline{Z}$ to illustrate its partial

		$N =$	1	2	3	4
Approximation	$\sum_{j=1}^{N} \langle \boldsymbol{v}^{(j)}, \boldsymbol{f}_3^{(N)} \rangle$		-0.0003	1.1715	0.8012	1.8502
Error	$-(\boldsymbol{g}^{(N+1)}, \boldsymbol{u}^{(N+1)})$		1.0033	-0.1691	0.2007	-0.8625
	$\sum_{j=1}^{N} \langle \boldsymbol{v}^{(j)}, \boldsymbol{f}_3^{(N)} \rangle - (\boldsymbol{g}^{(N+1)}, \boldsymbol{u}^{(N+1)})$		1.0030	1.0023	1.0019	0.9877
Observation	$\partial_r \overline{Z}$		1.0030	1.0030	1.0030	1.0030

Table 1: The error $-(\boldsymbol{g}^{(N+1)}, \boldsymbol{u}_3^{(N+1)})$ in the estimation of $\partial_r \overline{Z}$ using a truncation of the cumulant hierarchy at order N. See appendix B for further details.
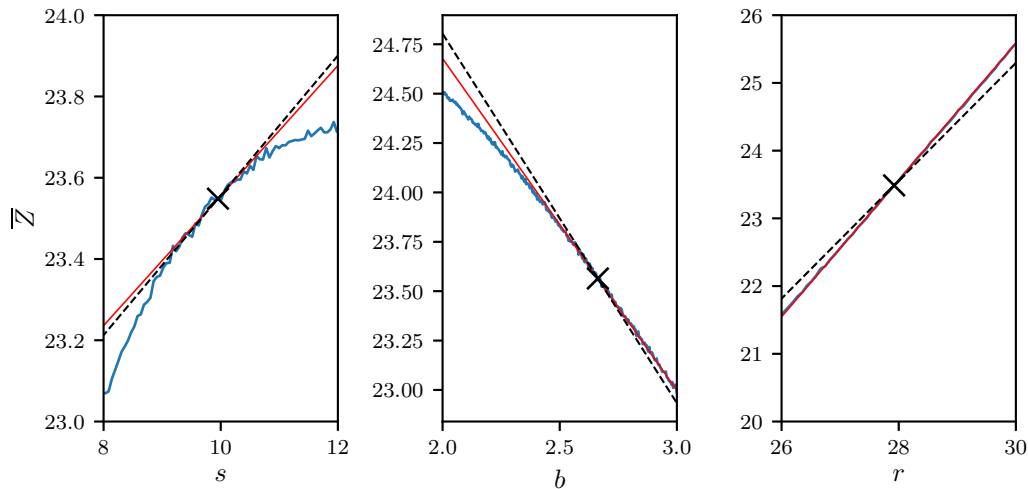
Figure 7: Local observation and approximation of the derivative of $\overline{Z}$ at the point $(s, b, r) = (10, 8/3, 28)$ (marked $\times$) with respect to the parameters s, b and r. The thick blue curve corresponds to observations from direct simulations of the Lorenz equations, in which one parameter is varied and the others are held constant. The straight red line corresponds to a local fit to the gradient by Wang (2013) and the dashed line corresponds to the gradient obtained from the optimised model approximation developed in §3.4.

	$\partial_s \overline{Z}$	$\partial_b \overline{Z}$	$\partial_r \overline{Z}$
Wang (2013, regression)	0.16	-1.68	1.01
Wang (2013)	0.21	-1.74	0.97
1st order	0.0000 (0.0000)	-8.8346 (-8.8327)	-0.0003 (-0.0001)
2nd order	0.0312 (0.0312)	1.5146 (1.5145)	1.1715 (1.1715)
3rd order	0.2144 (0.2145)	-2.7844 (-2.7840)	0.8012 (0.8012)
4th order	-0.3353 (-0.3350)	8.5854 (8.5774)	1.8502 (1.8495)
Model ($\tau = 10^3$)	0.2186	-2.5566	0.8379
Model ($\tau = 10^4$)	0.1754	-1.9172	0.8723
Model ($\tau = 10^5$)	0.1734	-1.9082	0.8730
Model ($\tau = 10^6$)	0.1712	-1.8743	0.8748

Table 2: Cumulant sensitivities for the Lorenz equations. Nth order corresponds to truncation of the cumulant equations at order N (i.e. discarding cumulants of order $N+1$, which is equivalent to assuming that $(\boldsymbol{u}^{(2)}, \boldsymbol{g}^{(2)}) = 0$), obtained from integrals over $\tau = 10^5$ dimensionless time units (values corresponding to $\tau = 10^6$ are shown in parentheses). The entries marked 'Model' correspond to those obtained by using an optimal eddy damping parameter μ_* in the equations for the third order cumulants.
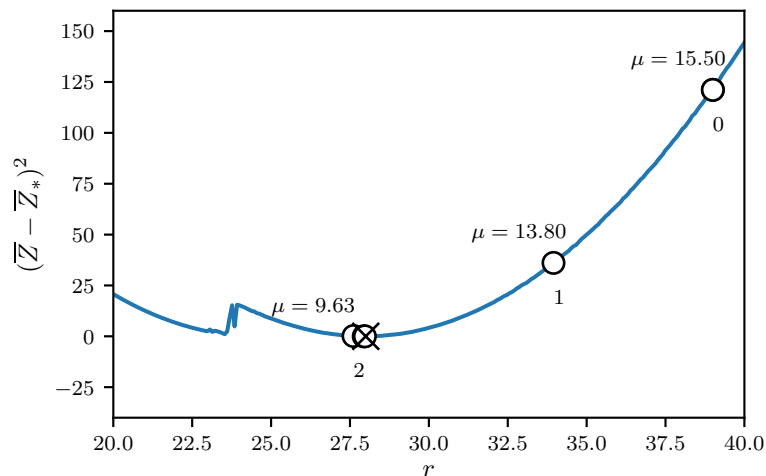
Figure 8: The solution to the inverse problem of determining the value of the renormalised Reynolds number r from a known observed statistic $\overline{Z}_*$. The circles denote points at which the functional and its gradients were evaluated to search for the minimum value $(\overline{Z} - \overline{Z}_*)^2$.

dependence on the parameters s, b and r. The gradients that are obtained by using the optimal model approach described in §3.4 are displayed in comparison with those that were obtained by linear regression analysis (Wang, 2013). The optimal model approach yields a reasonably good agreement with the observed gradients of $\overline{Z}$ at $(s, b, r) = (10, 8/3, 28)$. The optimal value of $\mu = \mu_*$ was found to be 8.96. A summary of the results, including the dependence of the computed gradients on the integration time used to obtain statistics, is provided in table 2.

4.4 *Test optimisation problem*

In practice, local gradient information can be used in a gradient-based optimisation routine. To demonstrate, we define the functional $(\overline{Z} - \overline{Z}_*)^2$ where $\overline{Z}_*$ corresponds to the desired value of $\overline{Z}$. For convenience we define $\overline{Z}_*$ as the value of $\overline{Z}$ corresponding to the parameters $(s, b, r) = (10, 8/3, 28)$ and attempt to solve the inverse problem of determining an *a priori* unknown r from the known value $\overline{Z}_*$. During each iteration of the optimisation procedure, we calculate the cumulants corresponding to a given set of parameters, and therefore the functional $\overline{Z}$. We then find the optimal eddy-damping parameter μ_*, before approximating the gradient of the functional. The optimal eddy-damping can therefore change at each step of the iteration procedure. We use the BFGS optimisation routine in the SciPy library and look for the parameter r_* corresponding to $\overline{Z}_*$. We set $s = 10$ and $b = 8/3$ and select $r = 39$ as an initial guess for r. Within four iterations the optimisation routine finds $r_* = 27.97$ and $\overline{Z}_*$ to within a tolerance of less than 0.001. This is an interesting, albeit contrived, example of a problem for which the use of sub-optimal gradients can nevertheless lead to an optimal solution because $\overline{Z} = \overline{Z}_*$ implies that $\nabla J = 0$, regardless of whether $\nabla \overline{Z} = 0$. More general optimisation problems, for which the value of an extremum might not be known in advance, will not necessarily share this property. Attempts to use the BFGS optimisation routine without providing local gradients were

unsuccessful.

5 Three-dimensional convection $(d = 9)$

A logical extension of the model for two-dimensional Boussinesq convection analysed in the previous section is the model for three-dimensional Boussinesq convection studied by Reiterer *et al.* (1998). Like its two-dimensional counter part, the system is a truncated Galerkin representation of the full dynamics. Unlike its two-dimensional counter part, the system evolves on a $d = 9$ dimensional, rather than $d = 3$ dimensional, phase space and therefore yields statistics that exhibit a more complicated dependence on the problem's parameters. Expressing temperature and velocity in terms of a triple Fourier series and retaining terms up to second order yields the following closed system of equations (Reiterer *et al.*, 1998):

$$\left.\begin{aligned}
\dot{Q}_0 &= -s\,b_1\,Q_0 - Q_1\,Q_3 + b_4\,Q_3^2 + b_3\,Q_2\,Q_4 - s\,b_2\,Q_6, \\
\dot{Q}_1 &= -s\,Q_1 + Q_0\,Q_3 - Q_1\,Q_4 + Q_3\,Q_4 - s\,Q_8/2, \\
\dot{Q}_2 &= -s\,b_1\,Q_2 + Q_1\,Q_3 - b_4\,Q_1^2 - b_3\,Q_0\,Q_4 + s\,b_2\,Q_7, \\
\dot{Q}_3 &= -s\,Q_3 - Q_1\,Q_2 - Q_1\,Q_4 + Q_3\,Q_4 + s\,Q_8/2, \\
\dot{Q}_4 &= -s\,b_5\,Q_4 + Q_1^2/2 - Q_3^2/2, \\
\dot{Q}_5 &= -b_6\,Q_5 + Q_1\,Q_8 - Q_3\,Q_8, \\
\dot{Q}_6 &= -b_1\,Q_6 - r\,Q_0 + 2\,Q_4\,Q_7 - Q_3\,Q_8, \\
\dot{Q}_7 &= -b_1\,Q_7 + r\,Q_2 - 2\,Q_4\,Q_6 + Q_1\,Q_8, \\
\dot{Q}_8 &= -Q_8 - r\,Q_1 + r\,Q_3 - 2\,Q_1\,Q_5 + 2\,Q_3\,Q_5 + Q_3\,Q_6 - Q_1\,Q_7,
\end{aligned}\right\} \tag{40}$$

where

$$\left.\begin{aligned}
b_1 &= 4\,\frac{1+k^2}{1+2\,k^2}, & b_2 &= \frac{1+2\,k^2}{2\,(1+k^2)}, & b_3 &= 2\,\frac{1-k^2}{1+k^2}, \\
b_4 &= \frac{k^2}{1+k^2}, & b_5 &= 8\,\frac{k^2}{1+2\,k^2}, & b_6 &= \frac{4}{1+2\,k^2}.
\end{aligned}\right\} \tag{41}$$

The parameters s and r continue to represent the Prandtl number and the renormalised Rayleigh number. In addition, equation (41) defines a set of geometrical parameters, as a function of the wave number k, which correspond to b in the previous problem. To within constants of proportionality, the variables X, Y and Z in the two-dimensional case correspond to Q_4, Q_9 and Q_6, respectively. More precisely, because $\overline{Z} = -A\,\overline{Q}_6$ for $A > 0$, we focus on the dependence of $-\overline{Q}_6$ on r, where $-Q_6$ is proportional to the strength of the horizontal average temperature with respect to a state of pure conduction. For details pertaining to the derivation of (40), the reader is referred to Reiterer *et al.* (1998). To aid comparison with the results presented in Reiterer *et al.* (1998), we choose $s = 10$, $k = 1/2$ and vary r. The statistics were obtained over a dimensionless time $\tau = 10^4$.

As described in Reiterer *et al.* (1998), when $r > 14.17$ for $s = 10$ and $k = 1/2$, the system is chaotic. When projected onto the Q_6, Q_9 plane the attractor consists of two lobes either side of the hyperplane $Q_9 = 0$, as can be seen in figure 9. As r increases the deviation of the horizontally averaged temperature from the linear behaviour associated with pure conduction increases.
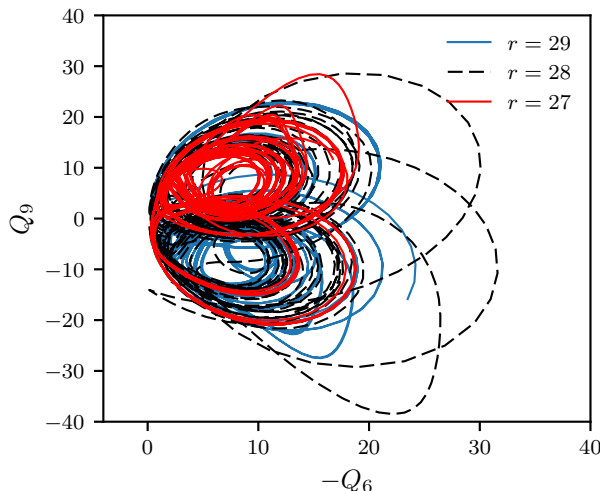
Figure 9: Projection of the chaotic attractor associated with the nine-dimensional system (40) onto the plane describing the mode associated with horizontally average temperature $-Q_6$ and the difference in temperature between ascending and descending fluid Q_9. The trajectory corresponds to 2000 dimensionless time units.

The precise relationship between r and $-Q_6$ for the parameters $s = 10$ and $k = 1/2$ is displayed in figure 10. In spite of the discontinuities resulting from the use of a finite time average for each value of r, the relationship indicates that $-Q_6$ tends to increase as r increases. At a glance, a linear relationship between r and $-Q_6$ over $[26, 30] \ni r$ appears to provide a reasonable first description of the sensitivity. However, closer inspection reveals that $-\partial_r \overline{Q}_6$ varies significantly on scales of approximately $\Delta r \sim 0.5$, in contrast to the equivalent relationship for the Lorenz system (see figure 1), for which $\partial_r \overline{Z}$ is approximately constant over a large range of r.

Since the dynamical system has $d = 9$ degrees of freedom, the number of cumulants up to order N is given by equation (22):

$$\sum_{j=1}^{N} \binom{j + d - 1}{d - 1} = 9 + 45 + 165 + \ldots + \binom{N + d - 1}{d - 1}. \tag{42}$$

The derivative of $-\overline{Q}_6$ with respect to r was computed by truncating the cumulant equations at $N = 3$ and invoking the optimal eddy-damping closure described in §3.4. As is evident from figure 10, the computed gradients appear to under estimate the underlying exact gradients in general, but nevertheless provide a reasonably good approximation. As pointed out in §3.4, the least-squares optimal eddy-damping parameter yields an approximation to the gradient based on point-wise observations, rather than the best approximation to the gradient. It is therefore useful to consider the sensitivity of the computed gradient to changes in the eddy-damping parameter μ by calculating the derivative $-\partial_\mu \partial_r Q_6$. Figure 10 displays gradients corresponding to the optimal eddy damping parameter μ_*, along with lines whose gradients are $\pm \mu_* \partial_\mu \partial_r Q_6(\mu_*)$ to indicate the sensitivity of the results to changes in μ. It is interesting that at $r = 29$, we observe that $\partial_\mu \partial_r Q_6 = 0$,
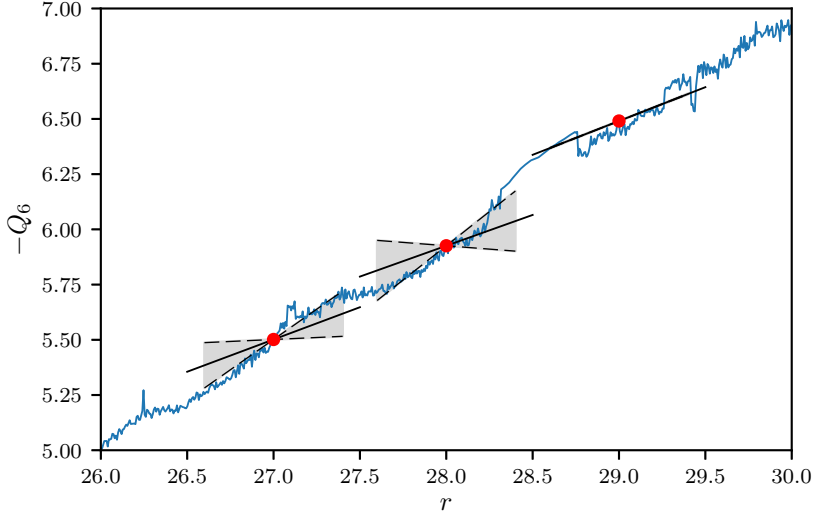
Figure 10: The dependence of $-Q_6$ on the system parameter r (blue line) and approximations to the local derivative using a truncated cumulant expansion. The statistics were obtained over a dimensionless time $\tau = 10^4$. The shaded regions indicate the linear sensitivity of the computed gradient to changes in the eddy-damping parameter μ; the gradient of the bounds of the shaded regions are calculated according to $\pm\mu_*\partial_\mu\partial_r Q_6$.

which indicates that the computed gradient is insensitive to changes in μ.

As discussed in §3.4, different values of μ correspond to different assumptions about the involvement of third-order cumulants in the statistical equilibrium. Picking an arbitrary value of μ in equation (38) might result in the adjoint operator being close to singular and therefore yielding gradients that depend sensitively on the choice of μ. To illustrate this, figure 11 shows evaluations of the derivative of $\overline{Q}_6$ with respect to r using equation (38) for values of μ in the vicinity of the optimal value μ_* as determined by equation (35). When $r = 27.0$ and $r = 28.0$ it is evident that some choices of μ result in a singular or near-singular adjoint operator and, therefore, a large amount of uncertainty in the resulting gradients. To obtain robust results in this particular case it is therefore necessary to use an optimal eddy-damping parameter that is determined systematically, rather than an estimation that is independent of observations. The optimal parameter μ_* appears to find a local maximum in the value of $-\partial_r\overline{Q}_6$ when $r = 29.0$, which explains why the estimated gradient is locally insensitive to changes in μ.

6   Conclusions

We have described a systematic means of obtaining approximate forward and adjoint sensitivity information from a chaotic system using a truncated system of cumulant equations. Unlike linearisation of the underlying evolution equations for individual trajectories, the cumulant equations yield robust, albeit approximate, information about functional derivatives. The method was designed for situations in which one has access to statistical data from the direct simulation of a potentially high-dimensional chaotic system and wishes
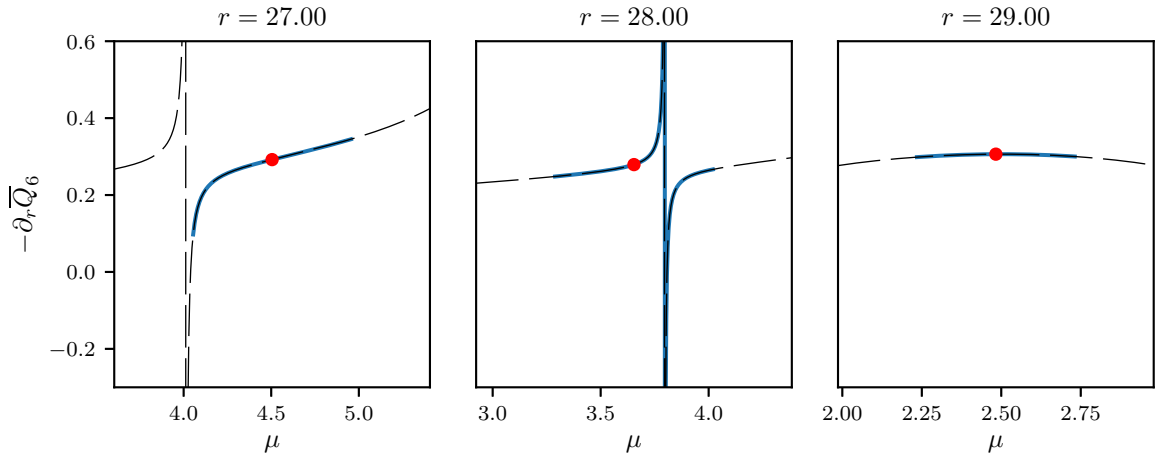
Figure 11: The dependence of the derivative $\partial_r \overline{Q}_6$, derived from a truncation of the cumulant hierarchy at $N = 3$, according to (38).

to approximate the gradients of a functional with respect to many input parameters. In principle the method could also be applied to obtain gradients of flow functionals from experimental measurements.

We combined data from direct simulation with tangent linear and adjoint equations for the system's statistical state dynamics. These equations can be obtained from the original system systemically using a cumulant generating function. Whilst the method is approximate, because it relies on truncation of the cumulant equations, the incorporation of observations to derive optimal truncations significantly improves its accuracy and robustness. Although the method itself is not restricted to statistically stationary problems, we expect the acquisition and incorporation of the corresponding unsteady statistical observations to be challenging.

The extraction of gradient information from functionals of chaotic dynamical systems is a stringent test for modelling and closure schemes. A given model can be tuned to adequately represent a given problem. However, unless it accurately describes the underlying physics, it is unlikely to yield accurate information about how an output functional changes with respect to changes in the problem's definition. Hence, the class of models from which one selects a suitable surrogate must be capable of describing the dynamics correctly. In the absence of shear, Rayleigh Bénard convection and, specifically, the Lorenz model, provide a difficult test for cumulant closures because truncation of the equations at second order removes interactions that are vital in determining the response of the system's statistical equilibrium. In contrast, for problems dominated by mean shear, such as jets, it is likely that cumulant truncation at second order would adequately capture the leading-order dynamics and would significantly simplify the approach to obtaining gradient information. The basic approach that we have described can be refined by exploring more appropriate ways of fitting the surrogate model.

Although we have focused on relatively low-dimensional dynamical systems, the idea of using cumulant expansions was motivated by the need to analyse high-dimensional dynamical systems. The challenge in the successful application of the method to large

systems lies in the acquisition of a large number of accurate high order cumulants and the systematic derivation and manipulation of a potentially large number of cumulant equations. In such cases statistical symmetries of a given problem can be used to significantly reduce the number of unknowns. An alternative or complementary approach would be to map the full system onto a relatively low-order model, for which the cumulants and their dynamics can be more readily obtained. Guided by the classical moment problem, further work should also incorporate restrictions that could be imposed on the gradients of cumulants to ensure that they point in a realisable direction.

A   Derivation of the cumulant equations

The Hopf generating functional (Hopf, 1952) is defined according to

$$\Psi(\boldsymbol{Q}(t), \boldsymbol{P}) = \exp\left(\imath P_i Q_i(t)\right), \tag{43}$$

where $\imath = \sqrt{-1}$. The moment $\overline{\boldsymbol{Q}^\alpha}$ can therefore be generated as

$$\overline{\boldsymbol{Q}^\alpha} = (-\imath)^{|\alpha|} \partial^\alpha \overline{\Psi}\Big|_{\boldsymbol{P}=\boldsymbol{0}}, \tag{44}$$

where $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_d)$ is a multi-index, such that $\boldsymbol{Q}^\alpha = Q_1^{\alpha_1} Q_2^{\alpha_2} \ldots Q_d^{\alpha_d}$ and $\partial^\alpha = \partial_{P_1}^{\alpha_1} \partial_{P_2}^{\alpha_2} \ldots \partial_{P_d}^{\alpha_d}$. A moment $\overline{\boldsymbol{Q}^\alpha}$ can be decomposed into a sum of products of cumulants U_β, containing all possible factorisations of the monomial $\boldsymbol{Q}^\alpha$:

$$\overline{\boldsymbol{Q}^\alpha} = \sum_{\pi \in \Pi(\alpha)} \prod_{\beta \in \pi} U_\beta, \tag{45}$$

where π is a multiset that decomposes a multi-index into addends. For example, if $\alpha = (2,1,0,0,\ldots)$ then $\pi = \{(2,0,\ldots),(0,1,0,\ldots)\}$ would be one such decomposition. The multiset $\Pi(\alpha)$ consists of all such decompositions. For example, if $\alpha = (0,0,4)$, then $\boldsymbol{Q}^\alpha = Z^4$, and

$$\begin{aligned}
\Pi(\alpha) = \{&\{(0,0,4)\}, \\
&\{(0,0,3),(0,0,1)\}^4, \\
&\{(0,0,2),(0,0,2)\}^3, \\
&\{(0,0,2),(0,0,1),(0,0,1)\}^6, \\
&\{(0,0,1),(0,0,1),(0,0,1),(0,0,1)\}\},
\end{aligned} \tag{46}$$

in which the exponents denote set multiplicities. In the example above, the set multiplicities arise from the different ways that a set consisting of 4 elements can be partitioned. According to (45) and (46), the moment $\overline{Z^4}$ can be expressed in terms of cumulants as

$$\overline{Z^4} = U_{004} + 4U_{003}U_{001} + 3U_{002}^2 + 6U_{002}U_{001}^2 + U_{001}^4. \tag{47}$$

The decomposition (45) is identical to that which arises when partial derivatives are applied to composite functions. Indeed, using $\Psi = \exp(\log(\Psi))$,

$$\partial^\alpha \overline{\Psi}\Big|_{\boldsymbol{P}=\boldsymbol{0}} = \sum_{\pi \in \Pi(\alpha)} \prod_{\beta \in \pi} \partial^\beta \log \overline{\Psi}\Big|_{\boldsymbol{P}=\boldsymbol{0}}, \tag{48}$$

which shows the logarithm of the moment generating function is the cumulant generating function.

B   Observed cumulant gradients

The gradients used to compute the truncation errors displayed in figures 5-6 were obtained from simulations of the Lorenz equations for 256 values of r uniformly distributed over a unit interval centred on $r = 28$. An approximation of the partial derivative of non-zero cumulants up to order 5 was obtained by minimising the squared difference between the straight line $(\partial_r J|_{r=28})r + J(0)$ and the data, which are both displayed in figure 12. The resulting gradients are tabulated in tables 3-6.

α	$\boldsymbol{g}^{(2)}$	$\boldsymbol{u}_3^{(2)}$
(2, 0, 0)	0.00e+00	2.68e+00
(1, 1, 0)	-3.75e-01	2.67e+00
(0, 2, 0)	0.00e+00	3.95e+00
(0, 0, 2)	0.00e+00	2.92e+00

Table 3: The constituent parts of the error $(\boldsymbol{g}^{(2)}, \boldsymbol{u}_3^{(2)})$ for truncation at order $N = 1$.

α	$\boldsymbol{g}^{(3)}$	$\boldsymbol{u}_3^{(3)}$
(2, 0, 1)	1.70e-03	2.19e+01
(1, 1, 1)	1.70e-02	7.79e+00
(0, 2, 1)	0.00e+00	1.14e+00
(0, 0, 3)	0.00e+00	-1.81e+00

Table 4: The constituent parts of the error $(\boldsymbol{g}^{(3)}, \boldsymbol{u}_3^{(3)})$ for truncation at order $N = 2$.

α	$\boldsymbol{g}^{(4)}$	$\boldsymbol{u}_3^{(4)}$
(4, 0, 0)	0.00e+00	-3.33e+02
(3, 1, 0)	2.56e-04	-3.33e+02
(2, 2, 0)	1.27e-04	-2.87e+02
(2, 0, 2)	-1.27e-04	1.33e+02
(1, 3, 0)	2.73e-04	-2.38e+02
(1, 1, 2)	-5.46e-04	-5.72e+00
(0, 4, 0)	0.00e+00	-2.33e+02
(0, 2, 2)	0.00e+00	-6.99e+01
(0, 0, 4)	0.00e+00	-2.32e+02

Table 5: The constituent parts of the error $(\boldsymbol{g}^{(4)}, \boldsymbol{u}_3^{(4)})$ for truncation at order $N = 3$.
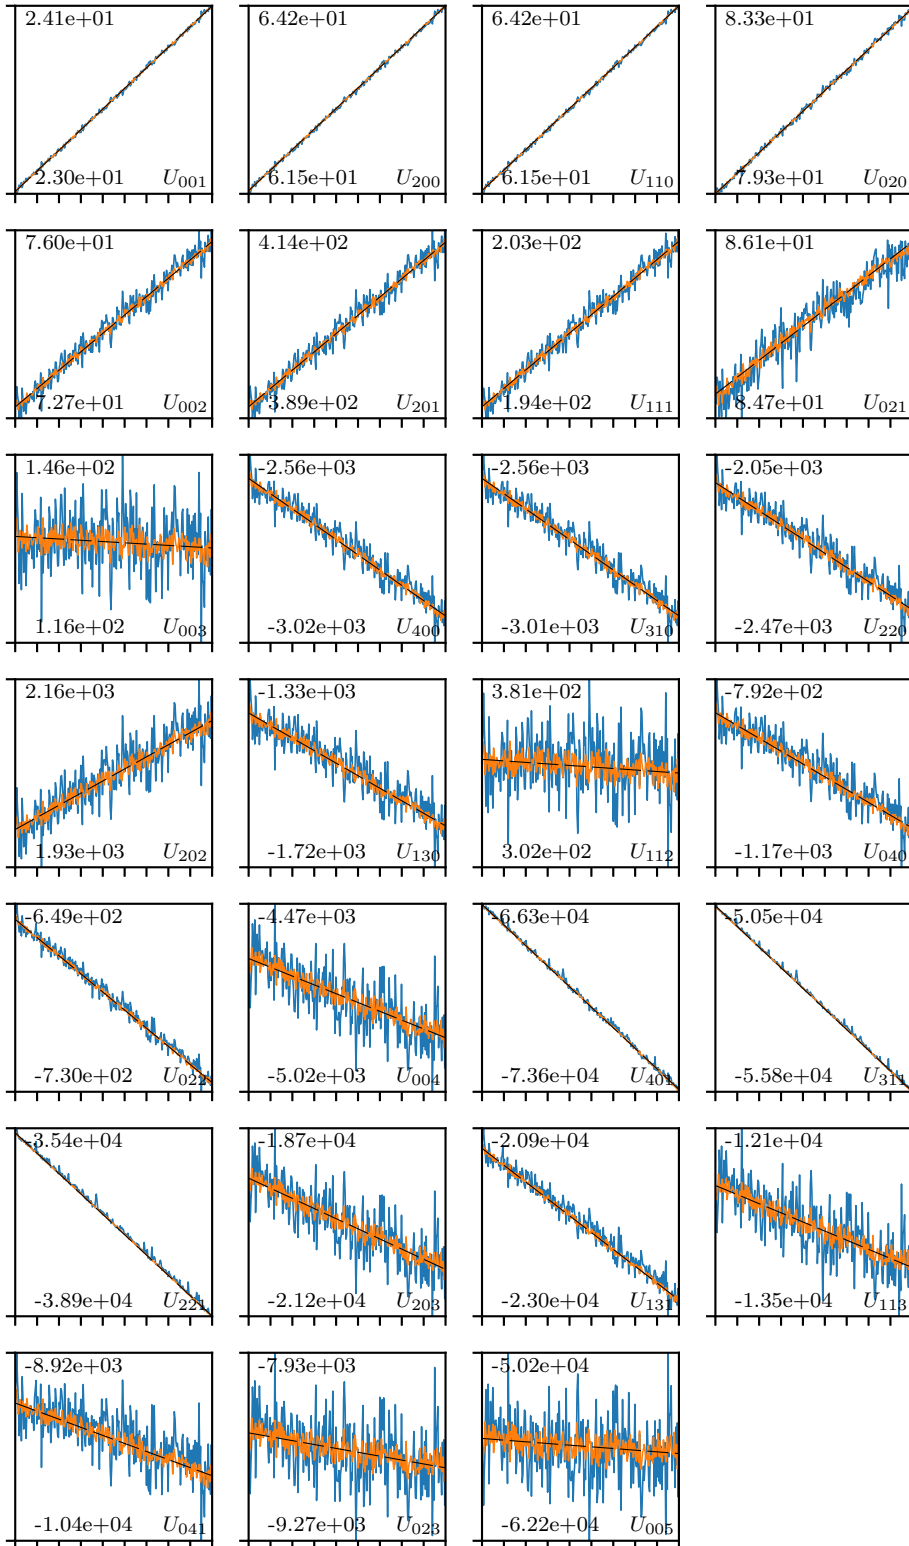
Figure 12: Estimators for the non-zero cumulants of the Lorenz attractor from simulations of duration $\tau = 10^5$ time units (blue/dark) and $\tau = 10^6$ time units (red/light). The gradients of the data were computed from simulations of duration $\tau = 10^6$ time units. The cumulant U_α corresponds to the moment $\overline{Q}^\alpha$ modulo all combinations of the corresponding low-order cumulants, as described in appendix A.

α	$\boldsymbol{g}^{(5)}$	$\boldsymbol{u}_3^{(5)}$
(4, 0, 1)	8.20e-06	-7.16e+03
(3, 1, 1)	-1.47e-04	-5.22e+03
(2, 2, 1)	-6.37e-06	-3.37e+03
(2, 0, 3)	-9.74e-06	-1.18e+03
(1, 3, 1)	-5.97e-05	-1.73e+03
(1, 1, 3)	-2.66e-05	-6.31e+02
(0, 4, 1)	0.00e+00	-5.73e+02
(0, 2, 3)	0.00e+00	-2.49e+02
(0, 0, 5)	0.00e+00	-9.62e+02

Table 6: The constituent parts of the error $(\boldsymbol{g}^{(5)}, \boldsymbol{u}_3^{(5)})$ for truncation at order $N = 4$.

Acknowledgements

References

Ait-Chaalal, F., Schneider, T., Meyer, B. & Marston, J. B. 2016 Cumulant expansions for atmospheric flows. *New Journal of Physics* .

Allawala, A. & Marston, J. B. 2016 Statistics of the stochastically forced Lorenz attractor by the Fokker-Planck equation and cumulant expansions. *Phys. Rev. E* **94**, 052218.

Auerbach, Ditza, Cvitanović, Predrag, Eckmann, Jean-Pierre, Gunaratne, Gemunu & Procaccia, Itamar 1987 Exploring chaotic motion through periodic orbits. *Phys. Rev. Lett.* **58**, 2387–2389.

Blonigan, P. J. & Wang, Q. 2014 Least squares shadowing sensitivity analysis of a modified Kuramoto-Sivashinsky equation. *Chaos, Solitons & Fractals* **64**, 16–25, nonequilibrium Statistical Mechanics: Fluctuations and Response.

Bohr, T., Jensen, M.H., Paladin, G. & Vulpiani, A. 2005 *Dynamical Systems Approach to Turbulence*. Cambridge University Press.

Cacuci, D.G. 2003 *Sensitivity & Uncertainty Analysis, Volume 1: Theory*. CRC Press.

Cooper, F. C. & Haynes, P. H. 2011 Climate sensitivity via a nonparametric fluctuation-dissipation theorem. *Journal of the Atmospheric Sciences* **68** (5), 937–953.

DIMET, F. LE & TALAGRAND, O. 1986 Variational algorithms for analysis and assimilation of meteorological observations: theoretical aspects. *Tellus A: Dynamic Meteorology and Oceanography* **38** (2), 97–110.

ECKMANN, J.-P. & RUELLE, D. 2004 *Ergodic theory of chaos and strange attractors*, pp. 273–312. New York, NY: Springer New York.

EYINK, G. L., HAINE, T. W. N. & LEA, D. J. 2004 Ruelle's linear response formula, ensemble adjoint schemes and Lévy flights. *Nonlinearity* **17** (5), 1867.

FARRELL, B. F. & IOANNOU, P. J. 2014 Statistical state dynamics: a new perspective on turbulence in shear flow. *arXiv* **1412.8290v1**.

FARRELL, B. F., IOANNOU, P. J., JIMÉNEZ, J., CONSTANTINOU, N. C., LOZANO-DURÁN, A. & NIKOLAIDIS, M.-A. 2016 A statistical state dynamics-based study of the structure and mechanism of large-scale motions in plane Poiseuille flow. *Journal of Fluid Mechanics* **809**, 290–315.

FARRELL, P. E., COTTER, C. J. & FUNKE, S. W. 2014 A framework for the automation of generalized stability theory. *SIAM Journal on Scientific Computing* **36** (1), C25–C48.

FOURES, D.P.G., CAULFIELD, C.P. & SCHMID, P.J. 2014 Optimal mixing in two-dimensional plane poiseuille flow at finite Péclet number. *Journal of Fluid Mechanics* **748**, 241–277.

FRISCH, U. 1995 *Turbulence: The Legacy of A. N. Kolmogorov*. Cambridge University Press.

GILES, M. B. & PIERCE, N. A. 2000 An introduction to the adjoint approach to design. *Flow, Turbulence and Combustion* .

HOPF, E. 1952 Statistical hydrodynamics and functional calculus. *Journal of Rational Mechanics and Analysis* **1**, 87–123.

JAMESON, A. 1988 Aerodynamic design via control theory. *Journal of Scientific Computing* **3** (3), 233–260.

KNOBLOCH, E. 1979 On the statistical dynamics of the Lorenz model. *Journal of Statistical Physics* **20** (6), 695–709.

KRAICHNAN, R. H. 1980 Realizability inequalities and closed moment equations. *Annals of the New York Academy of Sciences* **357** (1), 37–46.

LASAGNA, D. 2018 Sensitivity analysis of chaotic systems using unstable periodic orbits. *SIAM Journal on Applied Dynamical Systems* **17** (1), 547–580.

LEA, D., ALLEN, M. & HAINE, T. 2000 Sensitivity analysis of the climate of a chaotic system. *Tellus A* **52** (5).

LEITH, C. E. & KRAICHNAN, R. H. 1972 Predictability of turbulent flows. *Journal of the Atmospheric Sciences* **29** (6), 1041–1058.

LIONS, J.L. 1971 *Optimal control of systems governed by partial differential equations*. Springer-Verlag.

LORENZ, E.N. 1967 *The Nature and Theory of the General Circulation of the Atmosphere*. World Meteorological Organization.

LORENZ, E. N. 1963 Deterministic nonperiodic flow. *J. Atmos. Sci.* **20** (2), 130–141.

LUCAS, DAN & CAULFIELD, C. P. 2017 Irreversible mixing by unstable periodic orbits in buoyancy dominated stratified turbulence. *Journal of Fluid Mechanics* **832**.

LUCHINI, PAOLO & BOTTARO, ALESSANDRO 2014 Adjoint equations in stability analysis. *Annual Review of Fluid Mechanics* **46** (1), 493–517.

LUKACS, EUGENE 1970 *Characteristic functions*, 2nd edn. London: Griffin.

MARCHUK, G.I. 1995 *Adjoint Equations and Analysis of Complex Systems*. Springer.

MARCONI, U. M. B., PUGLISI, A., RONDONI, L. & VULPIANI, A. 2008 Fluctuation-dissipation: Response theory in statistical physics. *Physics Reports* **461**, 111–195.

MARSTON, J. B. & CONOVER, E. 2008 Statistics of an unstable barotropic jet from a cumulant expansion. *Journal of the Atmospheric Sciences* **65** (6), 1955–1966.

OGURA, Y. & PHILLIPS, N. A. 1962 Scale analysis of deep and shallow convection in the atmosphere. *J. Atmos. Sci.* **19** (2), 173–179.

PIRONNEAU, O. 1974 On optimum design in fluid mechanics. *Journal of Fluid Mechanics* **64**, 97–110.

REITERER, P., LAINSCSEK, C., SCHÜRRER, F., LETELLIER, C. & MAQUET, J. 1998 A nine-dimensional Lorenz system to study high-dimensional chaos. *Journal of Physics A: Mathematical and General* **31** (34), 7121.

ROTHMAYER, A. P. & BLACK, D. W. 1993 Ensembles of the Lorenz attractor. *Proceedings: Mathematical and Physical Sciences* **441** (1912), 291–312.

RUELLE, D. 2009 A review of linear response theory for general differentiable dynamical systems. *Nonlinearity* **22**, 855–870.

RUSSO, S. & LUCHINI, P. 2016 The linear response of turbulent flow to a volume force: comparison between eddy-viscosity model and DNS. *Journal of Fluid Mechanics* **790**, 104–127.

SEWELL, M. J. 1987 *Maximum and Minimum Principles*. Cambridge University Press.

SMALE, S. 1967 Differentiable dynamical systems. *Bulletin of the American Mathematical Society* pp. 747–817.

THUBURN, J. 2005 Climate sensitivities via a Fokker-Planck adjoint approach. *Quarterly Journal of the Royal Meteorological Society* **131** (605), 73–92.

Tobias, S. M., Dagon, K. & Marston, J. B. 2011 Astrophysical fluid dynamics via direct statistical simulation. *The Astrophysical Journal* **727** (2), 127.

Tobias, S. M. & Marston, J. B. 2013 Direct statistical simulation of out-of-equilibrium jets. *Phys. Rev. Lett.* **110**, 104502.

Vishnampet, R., Bodony, D. J. & Freund, J. B. 2015 A practical discrete-adjoint method for high-fidelity compressible turbulence simulations. *Journal of Computational Physics* **285**, 173–192.

Wang, Q. 2013 Forward and adjoint sensitivity computation of chaotic dynamical systems. *Journal of Computational Physics* **235**, 1–13.

Wang, Q., Hu, R. & Blonigan, P. 2014 Least squares shadowing sensitivity analysis of chaotic limit cycle oscillations. *Journal of Computational Physics* **267**, 210–224.

Yorke, James A. & Yorke, Ellen D. 1979 Metastable chaos: The transition to sustained chaotic behavior in the lorenz model. *Journal of Statistical Physics* **21** (3), 263–277.