



Approximating the linear response of physical chaos

Adam A. Śliwiak · Qiqi Wang

Received: 24 May 2022 / Accepted: 9 September 2022 / Published online: 27 September 2022
© The Author(s) 2022

Abstract Parametric derivatives of statistics are highly desired quantities in prediction, design optimization and uncertainty quantification. In the presence of chaos, the rigorous computation of these quantities is certainly possible, but mathematically complicated and computationally expensive. Based on Ruelle’s formalism, this paper shows that the sophisticated linear response algorithm can be dramatically simplified in higher-dimensional systems featuring statistical homogeneity in the physical space. We argue that the contribution of the SRB (Sinai–Ruelle–Bowen) measure gradient, which is an integral yet the most cumbersome part of the full algorithm, is negligible if the objective function is appropriately *aligned* with unstable manifolds. This abstract condition could potentially be satisfied by a vast family of real-world chaotic systems, regardless of the physical meaning and mathematical form

of the objective function and perturbed parameter. We demonstrate several numerical examples that support these conclusions and that present the use and performance of a simplified linear response algorithm. In the numerical experiments, we consider physical models described by differential equations, including Lorenz 96 and Kuramoto–Sivashinsky.

Keywords Chaos · Ruelle’s linear response theory · Sensitivity analysis · Space-split sensitivity (S3) · SRB measure gradient

1 Introduction

Linear response theory (LRT) [12] provides an array of mathematical methods for analysis of system’s reaction to small perturbations of imposed forces or control parameters. In particular, the linear response of a dynamical system should be understood as the derivative of its output with respect to an input parameter. The name “linear response” is a direct consequence of the Taylor series expansion, which indicates that the system’s reaction can be approximated by a linear function involving two terms: the unperturbed term and parametric derivative re-scaled by the imposed perturbation. Indeed, the use of Taylor series reveals one fundamental aspect of LRT. Namely, based only on information about the system in the unperturbed state, its response can be predicted for any small perturbation. Consequently, LRT is applicable to systems that vary

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s11071-022-07885-7>.

A. A. Śliwiak (✉) · Q. Wang
Center for Computational Science and Engineering,
Massachusetts Institute of Technology, 77 Massachusetts
Ave, Cambridge 02139, MA, USA
e-mail: asliwiak@mit.edu

A. A. Śliwiak · Q. Wang
Department of Aeronautics and Astronautics,
Massachusetts Institute of Technology, 77 Massachusetts
Ave, Cambridge 02139, MA, USA

Q. Wang
e-mail: qiqi@mit.edu

differentiably with respect to its input. Efficient numerical algorithms for approximating the linear response are fundamental in design optimization, uncertainty quantification, control engineering and inverse problems. These LRT-based computational tools are used in several fields of physics: electromagnetism [18], plasma physics and fusion [16], statistical physics [26], turbulent flows [23], climate dynamics [34], and many more.

In the presence of chaos, the classical formulation of LRT is modified. The reaction of a chaotic system is measured in terms of certain statistical quantities, e.g., long-time averages. Under the assumption of *ergodicity*, the statistics do not depend on initial conditions. Therefore, for a given chaotic model, the long-time statistics can be manipulated only by varying the input parameters. A prominent result in the field of LRT is the work of Ruelle [35,36], who rigorously derived a closed-form expression for the linear response of chaos. The major assumption of Ruelle's derivation is *uniform hyperbolicity*, which is a mathematical idealization of chaotic behavior. We postpone the description and explanation of this property for the following section of the paper. Solid numerical evidence found in the literature clearly indicates that uniform hyperbolicity is a sufficient, but not necessary, condition for the differentiability of statistics [5,7]. Indeed, these empirical results are consistent with *hyperbolic hypothesis* of Galavotti and Cohen [14]. This hypothesis presumes that several high-dimensional chaotic systems behave as though they were uniformly hyperbolic. It does not mean, however, that all properties of uniform hyperbolicity are satisfied by those systems, but several consequences following from this fundamental assumption could still be valid. This was clearly demonstrated in [28], where the author argued that the long-time averages computed for a 3D turbulence model are smooth despite local non-hyperbolic behavior.

While Ruelle's theory is regarded as one of the cornerstones in the field, its original expression for the linear response is impractical due to the *butterfly effect*, i.e., exponential growth of tangent solutions in time. The ensemble method proposed in [10] circumvented this problem by computing ergodic averages along several truncated trajectories. Despite its simplicity, that approach suffers from prohibitive computational costs induced by large variances of partial sensitivities. Shadowing methods [31,47] depart from the direct evaluation of Ruelle's expression by approx-

imating the shadowing trajectories [33], which lie in close proximity to the original orbit for a long period of time. Methods of this type have successfully been applied to high-dimensional fluid mechanics systems [5,28]. However, a recent study [8] demonstrated that shadowing trajectories may be *nonphysical* and that their statistical behavior could be dramatically different than that of the reference trajectory. This unwanted behavior had also been observed in earlier studies, e.g., in [5], which demonstrated large errors in shadowing-based approximations in spite of the apparently smooth behavior of the statistics. To the best of our knowledge, no rigorous studies that quantify or bound shadowing errors due to the problem of *nonphysicality* are available. An alternative way of computing the linear response involves the fluctuation–dissipation theorem (FDT) [19], which provides a time-convolution expression for the parametric derivative of statistics. FDT-based methods, such as the blended algorithm [1], require some physics-informed assumptions to accurately reconstruct the linear response operator.

Recent algorithmic developments rely on the regularized variant of Ruelle's expression. Indeed, as originally proposed by Ruelle in [35], one can apply integration by parts to the original formula in order to eliminate the product of Jacobians whose norm grows exponentially fast. However, since that formula involves Lebesgue integrals with respect to the Sinai–Ruelle–Bowen (SRB) measure [49] that is absolutely continuous only on unstable manifolds, an extra step is required before partial integration is applied. Namely, the input perturbation should be decomposed into two terms arranged in line with unstable and stable manifolds of the underlying dynamical system [35]. In the case of flows (continuous-time systems), the center manifold should also be taken into account in the perturbation splitting [37]. Based on this idea of regularization of Ruelle's closed-form expression, two conceptually similar methods for the linear response emerged in the past two years. Those are the fast linear response algorithm [29] and space-split sensitivity (S3) algorithm [9,40]. Neither of them introduces engineered approximations except for the ergodic-averaging required for the evaluation of Lebesgue integrals inherited from the original formula. They rigorously converge as a typical Monte Carlo procedure for any uniformly hyperbolic system. Methods of this type can be summarized as follows. Split linear response into two terms (or three terms if considering a flow), such that one uses

solutions of a regularized tangent equation (immune to the butterfly effect), while the second term requires computing the divergence on unstable manifolds. The unstable divergence directly follows from the partial integration on the expansive tangent subspace. One of the by-products is the SRB density gradient representing the divergence of SRB measure. This quantity is obtained by differentiating the measure preservation law, which effectively requires solving a series of regularized second-order tangent equations [29,39,41]. Differentiation of SRB measures, either explicit or implicit, is by far the most complicated and expensive part of both algorithms.

In this paper, we investigate whether and under what circumstances the complex numerical procedures for the linear response could be simplified. In particular, we attempt to answer the fundamental question about the significance of the SRB measure change. Rich numerical evidence found in the literature suggests that the computation of the SRB density gradient is not necessary to accurately approximate the linear response in a number of popular physical systems. For example, the aforementioned shadowing methods, which in fact regularize the tangent equation and do not compute the curvature of unstable manifolds, have been proven successful in 3D turbulence models [5,28]. Moreover, a recent theoretical study in [30] concludes that if both the input perturbation and objective function follow the multivariate normal distribution, the effect of the measure change is expected to decay proportionally to $\sqrt{m/n}$, where m is the number of positive Lyapunov exponents (LEs), while n denotes the system's dimension. That work, however, does not provide any numerical examples. Here, we show that the contribution of the unstable divergence could potentially be negligible if the objective function is specifically *aligned* with the unstable manifold. The meaning of *alignment* in this context is rigorously explained later in this work. Our numerical examples indicate that it is not uncommon that the SRB measure change is large and even has infinite variance, while its contribution to the linear response might be negligible at the same time. This paradox may have huge implications for approximating sensitivities in large physical systems. The only obstacle is an additional requirement for the objective function, which typically has a concrete physical meaning. Our argument is based on the fact that a vast family of practicable systems are statistically homogeneous in physical space. They include popular models governing climate

dynamics [17], turbulence [23], population dynamics [46], and several other phenomena. For such systems, we have freedom in representing any spatially averaged objective function, which effectively increases the probability of its alignment with a tangent subspace.

Our reasoning also relies on the specific orthogonal representation of the perturbation splitting proposed and numerically tested in [40]. In particular, we use orthogonal Lyapunov vectors to represent unstable manifolds everywhere on the attractor. Although they provide limited information on the geometry of the tangent space, there are three major reasons we favor orthogonal basis vectors over their covariant counterparts (CLVs). First, when ordered consistently with the decreasing set of LEs, both the Lyapunov basis sets have the same linear span [4]. This cascade property was used in [40] to stabilize the stable contribution of the S3 algorithm, as it enables us to orthogonally project out the unstable, unstable-center, or unstable-center-stable component of a tangent solution in a recursive manner. We also highlight the fact that S3 does not need stable directions alone. Second, the SRB measure change computed in the direction corresponding to the largest LE tends to be statistically smaller, even by orders of magnitude, compared to the other orthogonal directions. SRB measure slopes computed along the consecutive orthogonal directions are strongly correlated with the Lyapunov spectrum. We numerically verify this property and show that, when combined with the concept of *alignment* of the objective function, it may have a huge impact in controlling the magnitude of the unstable contribution. Finally, orthogonal Lyapunov bases are computationally cheaper compared to CLVs, as they require only a forward tangent solver with step-by-step QR factorization.

The structure of this paper is the following. In Sect. 2, we thoroughly review the space-split sensitivity (S3) algorithm for the linear response with an emphasis on potential problems. Subsequently, in Sect. 3, we explain the concept of alignment of the objective function and analyze its major implications in the context of the unstable contribution. A numerical experiment demonstrating a negligible effect of SRB measure change is presented. In Sect. 4, we conjecture that the alignment constraint is not an obstacle for higher-dimensional systems with statistical homogeneity. Based on our analysis, we propose a reduced variant of the S3 method and apply it to approximate the linear response of the Lorenz 96 and Kuramoto–Sivashinsky models. Sec-

tion 5 concludes this paper. Appendices A and B provide further technical details of S3: algorithm mechanics, implementation and cost analysis.

2 Space-split sensitivity (S3) method for chaotic flows

The purpose of this section is twofold. First, we review the main results of the linear response theory, i.e., Ruelle’s closed-form expression and its computable realization, known as the space-split sensitivity. Second, we present an extension of S3 to general hyperbolic flows and critically analyze its properties and major implications in the context of higher-dimensional systems.

Throughout this paper, we consider a parameterized n -dimensional ergodic flow,

$$\frac{dx}{dt} = f(x; s), \quad x(0) = x_0, \tag{1}$$

with $m \geq 1$ positive Lyapunov exponents, where s is a real-valued scalar parameter. The value of m approximates the dimension of the unstable (expanding) subspace, while particular LE values indicate the rate of exponential expansion/contraction [3]. Due to the assumed ergodicity, the statistical behavior of the system does not depend on the initial condition x_0 .

For a given smooth objective function $J : M \rightarrow \mathbb{R}$, our ultimate goal is to approximate the parametric derivative of the long-time average of J , defined as

$$\frac{d\langle J \rangle}{ds} := \frac{d}{ds} \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T J(x(t; s)) dt, \tag{2}$$

where M denotes the n -dimensional manifold defined by Eq. 1. We assume J does not depend on s .

2.1 Ruelle’s formalism and S3

Under the assumption of *uniform hyperbolicity*, Ruelle derived a closed-form expression for the linear response. Before we review the formula itself, we first focus on the assumption. A chaotic system is uniformly hyperbolic if its tangent space can be split into three *invariant* subspaces: unstable, stable and neutral. The first one and second ones are spanned by expanding and contracting directions of the tangent space, and they correspond to positive and negative LEs, respectively. These two subspaces, respectively, involve all tangent vectors

that exponentially increase and decay in norm along a trajectory. In this paper, we focus on autonomous flows, and thus, the tangent space also involves a neutral subspace that is parallel to the flow vector f and corresponds to the zero LE. In certain cases, a PDE-related dynamical system may involve more than one zero LE. For example, consider the Kuramoto–Sivashinsky equation with periodic boundary conditions. In this case, the neutral subspace is geometrically represented by a two-dimensional manifold (surface) that is tangent to f and spatial derivative of the solution at every point on the attractor. The key aspect of hyperbolicity is that the three subspaces are clearly separated from each other, which means that the smallest angle between them is far from zero everywhere on the attractor. Hyperbolic systems are structurally stable and admit the SRB measure μ [49], which contains the statistical description of the dynamics.

Assuming the system defined by Eq. 1 is uniformly hyperbolic, Ruelle’s linear response formula applies and can be expressed as follows [35,36],

$$\frac{d\langle J \rangle}{ds} = \sum_{t=0}^{\infty} \int_M D(J \circ \varphi^t) \cdot \chi d\mu, \tag{3}$$

where $g \circ h := g(h)$, $\chi = \partial_s \varphi \circ \varphi^{-1}$, $\varphi^t = \varphi(\varphi^{t-1})$, $\varphi^0(x) = x$, while D denotes the gradient operator (first derivative) in phase space. The diffeomorphic map $\varphi : M \rightarrow M$ can be interpreted as a time integrator of Eq. 1. For example, using the second-order explicit Runge–Kutta method (midpoint rule) with step size Δt , φ is related to f through the following relation,

$$x_{k+1} = \varphi(x_k) = x_k + \Delta t f(x_k + \frac{\Delta t}{2} f(x_k)). \tag{4}$$

Since the system is assumed to be ergodic, the Lebesgue integral with respect to measure μ can be approximated as,

$$\begin{aligned} \int_M h(x) d\mu &= \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T h(x(t)) dt \\ &\approx \frac{1}{N} \sum_{k=0}^{N-1} h(x_k) \end{aligned} \tag{5}$$

for any observable $h \in L^1(\mu)$ and a sufficiently large sample size N . Thus, the right-hand side (RHS) of Eq. 3 could potentially be approximated by computing a sufficiently long trajectory, ergodic-averaging the integrand per Eq. 5, and truncating the infinite series. However, note that

$$D(J \circ \varphi^t) \cdot \chi = (DJ)_t \cdot (D\varphi)_{t-1} \dots D\varphi \chi. \tag{6}$$

$(DJ)_t$ denotes the phase-space gradient of J evaluated t time steps into the future. To facilitate the notation, we will drop the parentheses, i.e., $(DJ)_t := DJ_t$. Therefore, unless χ is orthogonal to the unstable subspaces, the norm of that product grows exponentially fast with t ,

$$\|D\varphi_{t-1} D\varphi_{t-2} \dots D\varphi \chi\| \sim \mathcal{O}(\exp(\lambda_1 t)) \tag{7}$$

with $\lambda_1 > 0$, which means the direct evaluation of the RHS of Eq. 3 is computationally infeasible. The rate of exponential growth is determined by the leading LE denoted by λ_1 . Indeed, due to the butterfly effect, the derivative of the composite function $J \circ \varphi^t$ is the most problematic aspect of Ruelle’s original expression. Moreover, integration by parts is prohibited in this case, because one would also need to differentiate the SRB measure μ in the direction of χ . In general, the measure is absolutely continuous only on the expanding subspace [49]. Therefore, integration by parts would be possible only if χ belongs to unstable manifolds everywhere in M , which is generally not the case.

Motivated by the work of Ruelle [35,36], the authors of [7,9] proposed a new method, called the space-split sensitivity (S3), which regularizes Ruelle’s series for systems with one-dimensional unstable subspaces ($m = 1$). Based on its extension to general hyperbolic maps in [40], we derive and describe a space-split approach for chaotic flows with unstable manifolds of arbitrary dimension ($m \geq 1$). The main idea of S3, proposed in the aforementioned studies, is to decompose the perturbation vector χ into three terms,

$$\chi = \chi^u + \chi^c + \chi^s = \left(\sum_{i=0}^m c^i q^i \right) + (c^0 f) + \left(\chi - \sum_{i=0}^m c^i q^i - c^0 f \right), \tag{8}$$

such that χ^u and χ^c strictly belong to the unstable and neutral/center subspaces, respectively. In this splitting, $c^i, i = 0, \dots, m$ are some scalars that are differentiable on the unstable subspace defined by a local orthonormal basis $q^i, i = 1, \dots, m$. From now on, the superscript shall indicate the index of an array’s component. This notation does not imply exponentiation, unless explicitly stated otherwise. There are two major benefits of the perturbation splitting defined by Eq. 8:

- the unstable part of the linear response, i.e., the one involving χ^u , can now be integrated by parts, because it involves directional derivatives only along unstable subspaces,

- we can always find $c^i, i = 0, \dots, m$ through orthogonal projection such that the stable part (the one involving χ^s) of the linear response can be approximated by solving a regularized tangent equation that is bounded in norm.

We begin from exploring the second benefit of the splitting. Using the chain rule, one can rigorously show that the linear response defined by Ruelle’s series equals the ergodic average of $DJ \cdot v$, where v is a solution to the inhomogeneous tangent equation with χ as the source term. Thus, by replacing χ with χ^s in Eq. 3, we conclude that

$$\sum_{t=0}^{\infty} \int_M D(J \circ \varphi^t) \cdot \chi^s d\mu = \int_M DJ \cdot v d\mu, \tag{9}$$

where

$$v_{k+1} = D\varphi_k v_k + \left(\chi_{k+1} - \sum_{i=0}^m c_{k+1}^i q_{k+1}^i - c_{k+1}^0 f_{k+1} \right). \tag{10}$$

The subscript notation indicates the time step, i.e., $f(x(k\Delta t)) := f_k$, assuming uniform time discretization. To solve Eq. 10, we need to project out the unstable component of v , otherwise its norm will grow exponentially in time at the rate proportional to the largest LE. Moreover, we should also project out the component tangent to the center manifold to eliminate the increase of sample variances, which we illustrate later in Sect. 2.3. Therefore, we enforce v to be orthogonal to the unstable-center subspace by imposing a set of $m + 1$ constraints at every point on the manifold. Let $r_{k+1} := D\varphi_k v_k + \chi_{k+1}$ and, therefore,

$$(f_{k+1} \cdot f_{k+1}) c_{k+1}^0 = f_{k+1} \cdot \left(r_{k+1} - \sum_{i=1}^m c_{k+1}^i q_{k+1}^i \right), \tag{11}$$

$$c_{k+1}^i = q_{k+1}^i \cdot \left(r_{k+1} - c_{k+1}^0 f_{k+1} \right), \tag{12}$$

$i = 1, \dots, m.$

Equation 11–12 defines a linear system with $m + 1$ equations and $m + 1$ unknowns ($c^i, i = 0, 1, \dots, m$). The system’s matrix involves an $m \times m$ identity block I , while its Schur complement can be expressed as follows:

$$S_{k+1} = I - \frac{Q_{k+1}^T f_{k+1} (Q_{k+1}^T f_{k+1})^T}{f_{k+1} \cdot f_{k+1}}, \tag{13}$$

where Q is an $n \times m$ matrix containing an orthonormal basis of the unstable manifold, $q^i, i = 1, \dots, m$. Thus,

the coefficients $c^i, i = 1, \dots, m$, stored in the array c , are obtained by solving the following reduced system,

$$S_{k+1} c_{k+1} = Q_{k+1}^T \left(r_{k+1} - \frac{f_{k+1} \cdot r_{k+1}}{f_{k+1} \cdot f_{k+1}} f_{k+1} \right), \tag{14}$$

while c^0 is computed directly from Eq. 11. We conclude that the stable part of the linear response can be evaluated through the ergodic average of $DJ \cdot v$ (see Eq. 5), where v satisfies Eq. 10–12.

The next step is the neutral contribution, which involves the perturbation component that is parallel to f . Analogously to Eq. 6, we can expand

$$\begin{aligned} D(J \circ \varphi^t) \cdot \chi^c &= D(J \circ \varphi^t) \cdot (c^0 f) \\ &= c^0 DJ_t \cdot (D\varphi_{t-1} \dots D\varphi f). \end{aligned} \tag{15}$$

Applying the Taylor series expansion, we note that

$$\begin{aligned} f(\varphi(x)) &= f(x) + Df(x) (\varphi(x) - x) \\ &\quad + \mathcal{O}((\varphi(x) - x)^2), \end{aligned} \tag{16}$$

and, analogously,

$$\varphi(x) = x + \Delta t Df(x) + \mathcal{O}(\Delta t^2). \tag{17}$$

By differentiating Eq. 17 and plugging it to Eq. 16, we notice that in the limit $\Delta t \rightarrow 0$ we retrieve the covariance property, which reads

$$f(\varphi(x)) = D\varphi(x) f(x). \tag{18}$$

This implies that the neutral part can be simplified to

$$\begin{aligned} \sum_{t=0}^{\infty} \int_M D(J \circ \varphi^t) \cdot \chi^c d\mu &= \sum_{t=0}^{\infty} \int_M c^0 DJ_t \cdot f_t d\mu \\ &= \sum_{t=0}^{\infty} \int_M c_{-t}^0 DJ \cdot f d\mu. \end{aligned} \tag{19}$$

Equation 19 means that the neutral part of the linear response equals the infinite series of k -time correlations between c^0 , which is computed for the stable part, and $DJ \cdot f$. Under the assumption of uniform hyperbolicity, for any two Hölder-continuous observables J and h , k -time correlations exponentially converge to the product of expected values as $t \rightarrow \infty$ [11, 49], i.e.,

$$\left| \int_M (J \circ \varphi^t) h d\mu - \int_M J d\mu \int_M h d\mu \right| \leq C\delta^t \tag{20}$$

for some $C > 0$ and $\delta \in (0, 1)$. In the context of the linear response theory, at least one of the observables has zero expectation with respect to μ . Using this property, we approximate the neutral part by truncating the infinite series and computing each Lebesgue integral through Eq. 5.

The final missing contribution of the total linear response is the unstable term. Indeed, this is the only term we can apply integration by parts to, which yields [40]

$$\begin{aligned} &\sum_{t=0}^{\infty} \int_M D(J \circ \varphi^t) \cdot \chi^u d\mu \\ &= \sum_{t=0}^{\infty} \sum_{i=0}^m \int_M c^i \partial_{q^i} (J \circ \varphi^t) d\mu \\ &= - \sum_{t=0}^{\infty} \sum_{i=1}^m \int_M (J \circ \varphi^t) (c^i g^i + b^{i,i}) d\mu, \end{aligned} \tag{21}$$

where

$$b^{i,j} := \partial_{q^j} c^i, \quad g^i := \frac{\partial_{q^i} \rho}{\rho}, \tag{22}$$

the operator $\partial_{q^i}(\cdot) := D(\cdot) \cdot q^i$ denotes the directional derivative along q^i in phase space, while ρ denotes the density of the SRB measure μ conditioned on an unstable manifold. Several intermediate steps are required to derive the RHS of Eq. 21. First, the SRB measure is disintegrated across parameterized unstable manifolds. Second, partial integration is applied within each parameterized subspace. The resulting boundary terms vanish as proven in [35], which implies that in all integral transformations of this type, the boundary integrals can be neglected. The reader is also referred to [41] for a detailed description of every step of this process and relevant numerical examples. The major implication of Eq. 21 is that the composite function $J \circ \varphi^t$ is no longer differentiated, but there are two new quantities that must be computed instead. A rigorously convergent recursive algorithm for b and g has recently been proposed in [40]. That algorithm requires solving a collection of first- and second-order tangent equations, and was developed for discrete chaotic systems. In Appendix A, we extend it to hyperbolic flows and analyze its cost. Notice that if g and b are available, then, analogously to the neutral part, the unstable term is expressed in terms of infinite series of k -time correlations.

To summarize, the space-split method regularizes Ruelle’s original expression by splitting it into three major parts: stable, neutral and unstable. Each of them can be approximated through ergodic-averaging of a single (in stable part) or many (in neutral and unstable parts) ingredients. Recent rigorous [9] and computational [40] studies have shown that the rate of convergence of all linear response parts is approximately

proportional to $1/\sqrt{N}$, where N denotes the trajectory length. We highlight the fact that these studies were restricted to hyperbolic systems only. Thus, the S3 method is in fact a Monte Carlo procedure that relies on recursive formulas in the form of tangent equations that are executed to find g, b, v and other necessary quantities.

2.2 Numerical example: Lorenz 63

To test the space-split algorithm (see Algorithm 2), we shall consider the three-dimensional Lorenz 63 system,

$$\begin{aligned} \frac{dx}{dt} &= \sigma(y - x), & \frac{dy}{dt} &= x(\rho - z) - y, \\ \frac{dz}{dt} &= xy - \beta z, \end{aligned} \tag{23}$$

which is one of the simplest chaotic flows. This ODE system models thermal convection of a fluid cell that is warmed from the one side and cooled from the opposite side. The original study of this model [24] demonstrated chaotic behavior at $\sigma = 10, \beta = 8/3, \rho \gtrsim 24$. For this choice of parameters, the strange attractor has a characteristic butterfly-shaped structure. The purpose of our experiment is to approximate the derivative of the long-time average of $J = J(z)$ with respect to the Rayleigh parameter ρ using S3. In this section, ρ should not be confused with the SRB measure density. Figure 1 illustrates the behavior of the statistics of two different objective functions, as well as the three Lyapunov exponents for $\rho \in [20, 40]$. We observe that λ_1 becomes positive for $\rho \gtrsim 24$, which is consistent with the original study. The presence of a zero LE indicates there exists a tangent subspace that is parallel to the flow, which is typical for autonomous chaos. Note that, in the chaotic regime, both long-time averages seem to be differentiable in the considered parametric space.

To integrate Eq. 23 in time, we used the second-order explicit Runge–Kutta with step size $\Delta t = 0.005$. As described in Appendix A, the space-split algorithm requires a few evaluations of first- and second-order differentiation operators of φ every time step. For this particular time integrator, the computation of $D^2\varphi(\cdot, \cdot)$ involves three evaluations of the Hessian of f , per our derivations in Appendix B. Fortunately, in the case of the Lorenz 63 system, $D^2f(\cdot, \cdot)$ is constant, which significantly reduces the cost.

The S3 algorithm relies on several recursive formulas in the form of tangent equations. Earlier studies




Fig. 1 Long-time averages of two different objective functions (upper) and Lyapunov exponents (lower) versus the Rayleigh parameter ρ . Ergodic averages have been taken over $N\Delta t = 50,000,000$ and $N\Delta t = 5,000$ time units, respectively

[9,40] proved both analytically and numerically that these recursions converge exponentially fast in discrete hyperbolic systems. We numerically investigate whether these results still apply to the Lorenz 63 flow. The upper plot of Fig. 2 illustrates a convergence test for three different quantities: SRB density gradient g , tangent solution v and its directional derivative (along q) w . These are three major ingredients that contribute to the total linear response. Along a single trajectory, we impose two different initial conditions for v, w and a (note $g = -q \cdot a$) and compute the norm/absolute value of the two solutions. The semi-logarithmic plot clearly indicates that all the norms decrease exponentially in time with a short transition at the beginning of simulation. To obtain a machine-precision approximation of

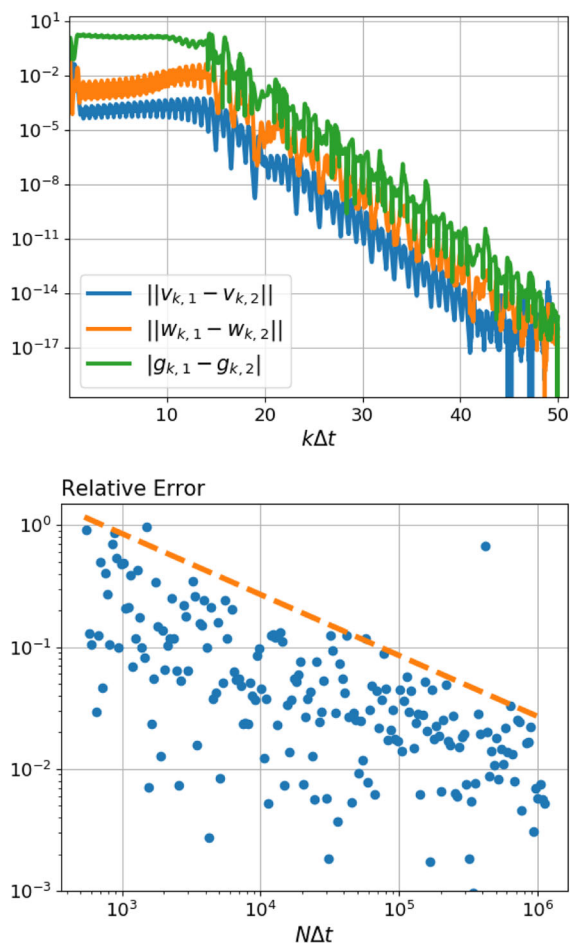


Fig. 2 Upper: Relation of the norm/absolute value of the difference between quantities obtained along two different random orbits, labelled as 1 and 2, and time-averaging window $k\Delta t$. Lower: Relative error of the linear response approximation versus time-averaging window, computed for $J = z$ at $\rho = 28$. 200 independent simulations were run at a logarithmically uniform grid of $N\Delta t$. The dashed line represents a function $C/\sqrt{N\Delta t}$, $C > 0$

these quantities, we need only 50 time units. A similar behavior has been observed in the case of discrete systems [40]. We use this result to set the truncation parameter to $T\Delta t = 100$ in our simulations to guarantee all ergodic-averaged quantities are very close to their true values. Another property of the S3 algorithm is the convergence rate of its final output, $\langle J \rangle / d\rho$, with respect to the time-averaging window $N\Delta t$. Indeed, a truncation of the trajectory by choosing a finite N is the only non-negligible source of error of the entire numerical procedure. The lower plot of Fig. 2 shows the decay of the relative error of the linear response

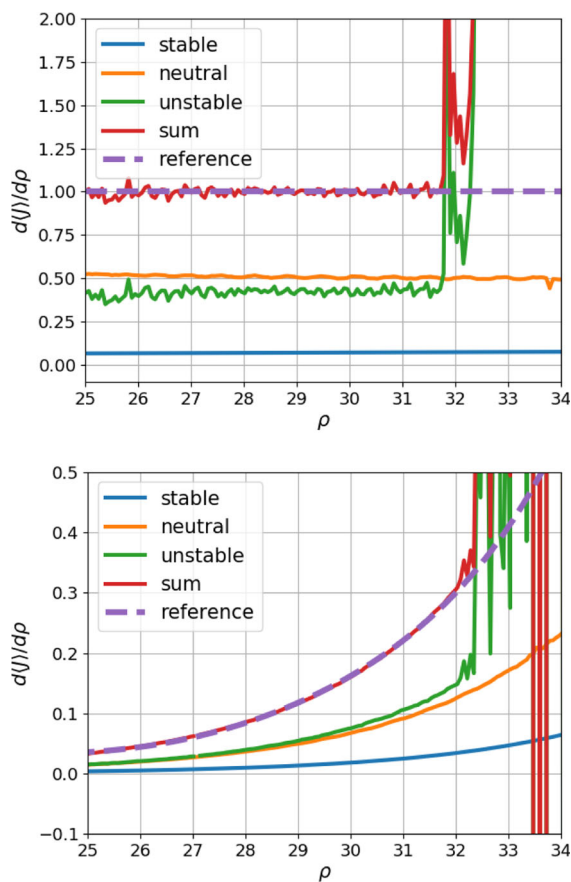


Fig. 3 Output of Algorithm 2 generated for $J = z$ (upper) and $J = \exp(x/4)/10000$ (lower) at 144 values of ρ distributed uniformly. Each simulation was run for $N\Delta t = 1,000,000$ time units. The reference solution (dashed curve) was obtained using central finite differences and data shown in Fig. 1. Before differentiation, we interpolated the data using first- and sixth-order polynomial fits, respectively

approximation, which is computed with respect to the finite difference approximation of the slope of statistics generated in Fig. 1. We observe that the error trend confirms theoretical predictions, which means that S3 behaves as a typical Monte Carlo simulation.

In our simulations, we truncate the infinite series by setting $K\Delta t = 50$, where K represents the number of series terms contributing to the numerical approximation. The optimal value of $K\Delta t$ should be relatively small, given the exponential decay of correlations. In [40], the reader will find a more detailed study about the impact of K on the error. Based on the convergence study and our discussion above, we run Algorithm 2 for Lorenz 63 ($n = 3, m = 1$) to compute parametric derivatives of the long-time averages illustrated in




Fig. 4 Discrete values of the stable integrand $DJ \cdot v$ computed using the S3 version described in Sect. 2.1 (red) and its “discrete” counterpart from [40] (blue). This simulation was performed for Lorenz 63 at $\rho = 28$. The solid lines represent the standard deviations of $DJ \cdot v$ collected from the beginning of the simulation until k th step. The dashed line represents a linear function. (Color figure online)

Fig. 1 at $\rho \in [25, 40]$. Figure 3 shows the behavior of the obtained linear response approximations. For a wide range of Rayleigh constant values, S3 provides accurate estimations of the sensitivities. Indeed, for $\rho \in [25, 32.3]$ we observe good agreement between the total sensitivity, denoted by “sum”, and corresponding reference values. At $\rho \approx 32.3$, the S3 approximation diverges due to the collapse of the unstable part. Note that, in both cases, the stable contribution is small compared to the two other terms. In the following section, we further explore the encountered problem and summarize critical aspects of the presented algorithm.

2.3 Critical view on S3

In the context of approximating linear response of higher-dimensional chaos, we shall investigate potential problems of the S3 algorithm. In particular, we focus on dynamical properties of chaotic flows that might lead to numerical difficulties. Some algorithmic challenges, including the computational cost, are also discussed.

2.3.1 Special treatment of the neutral component

In Sect. 2.1, we derived a numerical scheme based on the three-term linear splitting in Eq. 8. Indeed, there is

a subtle difference between this splitting and the one proposed for discrete systems. In the former, the neutral term is treated separately thanks to which the stable term includes only tangent solutions that are parallel to the unstable-center subspace. In Fig. 4, we plot discrete values of the stable integrand $DJ \cdot v$ obtained for Lorenz 63 at $\rho = 28$ using both versions of S3. We notice that if the neutral direction is not projected out from the tangent solution, then the standard deviation of $DJ \cdot v$ grows linearly with time. The extra projection against f guarantees the standard deviation is approximately constant.

While the convergence of the Monte Carlo procedure is now guaranteed, the extra projection requires assembling, inverting, and differentiating the Schur complement. As described in Appendix A, that minor conceptual adjustment requires major modifications of the “discrete” version of S3.

2.3.2 Problem with hyperbolicity and SRB measure gradient

Recall that the fundamental assumption of Ruelle’s formalism is hyperbolicity. Any form of linearly separated perturbation splitting that enables partial integration and that guarantees boundedness of the stable part, e.g., the one presented in this paper or the shadowing-based variant proposed in [29], is sufficient to construct stable numerical schemes. However, the dynamical structure of many chaotic flows, including the simple Lorenz 63 system, does not satisfy all basic properties of hyperbolicity.

In Fig. 5, we illustrate the distribution of tangency measures $0 \leq \alpha \leq 1$ between two pairs of subspaces: (1) unstable and center, (2) unstable-center and stable, along a random trajectory of Lorenz 63 at different values of the Rayleigh parameter. To generate these plots, we used the fast algorithm for hyperbolicity verification proposed by Kuptsov in [20]. The two measures we compute, respectively, represent d_1 , and $2 d_2$, which are rigorously defined by Eq. 7 in that work. The parameter α is closely related to the minimum angle between two subspaces normalized by $\pi/2$ as pointed out and tested in [44]. If the statistical distribution of α is not strictly separated from the origin, i.e., the corresponding PDF has nonzero values at $\alpha \approx 0$, then several tangencies of a given subspace pair are highly likely to occur. We observe that, regardless of the choice of ρ , there exist tangencies between the unstable and cen-




Fig. 5 Distribution of the normalized measure α between unstable/center subspaces (blue PDF) and unstable-center/stable subspaces (orange PDF). They have been computed along a random trajectory of the Lorenz 63 system for 5000 time units. To increase the accuracy of PDFs, we used the fourth-order Runge–Kutta time integrator with $\Delta t = 0.005$

ter subspaces. Several numerical examples presented in [20] imply that the absence of unstable-center separation is a common property of several physical systems. However, for some ρ , the Lorenz 63 system admits splitting of the tangent space into unstable-center and stable subspaces. This behavior has been known in the literature [27] under the name of *singular hyperbolicity*. Note that the Lorenz 63 oscillator loses this property at ρ between 30 and 35, which coincides with the collapse of the S3 algorithm. In particular, the unstable term blows up within this parameter regime, which indicates that μ becomes rough along expansive directions. From the study on differentiability of statistics of the Lorenz 63 system [43], we learn that the SRB density gradient g is Lebesgue integrable, i.e., $g \in L^1(\mu)$, only if $\rho < 32$. If ρ is close to the value of 28, then g is even square-integrable. The authors of the same paper argue that the integrability of g is both necessary and sufficient condition for differentiability of statistics. We conclude that even if Eq. 3 holds, one still needs to handle the by-products of partial integration, which might pose a serious challenge for Monte Carlo algorithms requiring pointwise values of derivatives of μ and other observables.

The smoothness of the SRB measure is not guaranteed in non-hyperbolic systems, which means that some components of g might not exist at all at some points on the attractor. Indeed, numerical experiments presented in [20, 44] indicate that some higher-dimensional physical systems, e.g., the Ginzburg–Landau equation, are clearly non-hyperbolic. Similar numerical results were provided for a 3D turbulent flow in [28]. Since g is an integral part of the S3 procedure and its value is computed everywhere along a random trajectory, we expect that the unstable contribution might blow up in the case of such systems.

2.3.3 Implementation and cost

We shall now comment on practical aspects of the full linear response algorithm, which is described in Appendix A. In terms of the implementation, both the stable and neutral parts do not require significant changes of the existing tangent/adjoint solvers. The former is obtained by solving a collection of first-order tangent equations. They are stabilized by step-by-step elimination of unstable-center tangent components through QR factorization that is needed to find a new basis of the subspace (matrix Q) and the Jaco-

bian of coordinate transformation (matrix R). The R factor can also be used to approximate m largest LEs, which is indeed a very useful by-product of the proposed algorithm [4]. The unstable contribution requires the implementation of the second-order derivative operator, which is necessary for g and b . While this is generally not a problem for simple systems, the need for a second-order tangent solver might require extra tools, such as automatic differentiation packages, for complicated higher-dimensional models.

It turns out that the presence of the Hessian is not the major burden of the full S3 algorithm. The typical structure of large physical systems is sparse due to the localized stencils of the most popular spatial discretization schemes. Therefore, the computational cost of matrix-vector or tensor-vector products is typically linear in n . Two other factors that determine the total cost are the trajectory length N and the number of positive LEs m . The former defines the accuracy of ergodic-averaging and indicates the number of primal and tangent solution updates and thus contributes linearly to the total cost. Based on our estimate in Appendix A, the final cost is proportional to the third power of m . The most expensive chunk of the algorithm is associated with the SRB density gradient g , which requires solving $\mathcal{O}(m^2)$ second-order tangent equations that is followed by a stabilizing normalization procedure consuming extra $\mathcal{O}(nm^3)$ flops. This might pose a serious challenge for systems with hundreds of unstable modes, such as 3D turbulence models.

2.3.4 Future prospects

The non-approximative methods for computing linear response of chaotic systems, such as the S3 algorithm, provide a rich collection of numerical tools for analysis of the underlying dynamics. Its major drawback is that the derivation of its components relies on the assumption of hyperbolicity and smooth SRB measure. These properties might be violated leading to the collapse of some parts of the full S3 algorithm. Nevertheless, we acknowledge the growing popularity and interest in hyperbolic systems among physicists and engineers. In a comprehensive review book of Kuznetsov [22], the author justifies this trend and provides several examples of hyperbolic attractors describing physical phenomena.

Despite the problems with hyperbolicity and large costs, can we still use some parts of the S3 algorithm

to find accurate estimates of linear response for higher-dimensional systems? As argued in [43], the collapse of the algorithm for g does not necessarily mean the linear response does not exist. Indeed, several aforementioned studies involving sensitivity analysis of large systems numerically demonstrate that their statistics are indeed differentiable. Figure 3 indicates that both the neutral and stable contributions of Lorenz 63 remain “stable” over the entire parametric regime. Removal of the unstable contribution would dramatically reduce the cost of S3, as the expensive and potentially incomputable g would no longer be needed. In the case of Lorenz 63, however, the unstable contribution accounts for approximately 40% of the total sensitivity. Therefore, omission of the unstable contribution of this system would give rise to significant errors. This observation leads to a fundamental question. Are there systems whose unstable contribution is small and can be neglected? If so, are they relevant for practitioners? We try to answer those questions in the remainder of this paper.

3 Unstable contribution: Can we neglect that term?

As we pointed out in Sect. 2.3, the computation of the unstable part of the linear response might be cumbersome due to several reasons. The purpose of this section is to provoke a discussion about the significance of that term. In particular, we shall present some evidence, indicating that the unstable term could be negligible and thus completely neglected if certain conditions are met.

Let us consider the leading term of Eq. 21, i.e., the one corresponding to $t = 0$,

$$U := \int_M J d d\mu, \quad d := d_{cg} + d_b := \sum_{i=1}^m c^i g^i + b^{i,i}. \tag{24}$$

Assuming the exponential decay of correlations holds, it is clear that the whole infinite series is small if U is small. Applying the Cauchy–Schwarz and triangle inequalities, we upper bound the magnitude of U ,

$$|U| \leq \|J\|_2 \|d\|_2 \leq \|J\|_2 (\|d_{cg}\|_2 + \|d_b\|_2), \tag{25}$$

where $\|\cdot\|_2$ denotes the L^2 norm with respect to μ defined as

$$\|h\|_2 := \sqrt{\int_M h^2 d\mu} \tag{26}$$

for any scalar function $h \in L^2(\mu)$. According to Inequality 25, we see that a small L^2 norm of the unstable divergence d implies that the entire unstable contribution is negligible as well. Recall that the vector c represents projections of the tangent solution v onto the unstable subspace, which depends on both χ , i.e., the parametric perturbation of the system, and geometry of the unstable manifold. The final term contributing to $\|d\|_2$ is the SRB density gradient, which represents measure change in m orthogonal directions of the unstable subspace. These directions, stored in the Q matrix, indicate how the unperturbed trajectory deforms in time. The rate of geometric expansion in the i -th direction is reflected by the i -th Lyapunov exponent λ_i , whose value can be expressed in terms of the following ergodic average [13],

$$\lambda_i = \int_M \log |q^i(\varphi(x)) \cdot D\varphi(x) q^i(x)| d\mu. \tag{27}$$

We also acknowledge that the computation of Q is an integral part of the S3 procedure (see Appendix A). In that algorithm, the columns of Q are sorted from the most expansive ($i = 1$) to the least expansive ($i = m$) direction. Equation 27 implies that a bunch of infinitesimally close points will scatter very fast along the q^i direction if λ_i is large, resulting in a small local measure change. In other words, larger expansion rates lead to the dilution of measure, which consequently decreases the measure gradient. Therefore, assuming the positive LEs are separated from each other, we conjecture that the measure change along q^1 and q^m is expected to be the smallest and largest, respectively. In particular, if $\lambda_1 > \lambda_2 \dots > \lambda_m$, then

$$\|g^1\|_2 < \|g^2\|_2 < \dots < \|g^m\|_2.$$

We verify this presumption later in a numerical experiment. Its major consequence is that we can potentially find two different directions on unstable manifolds along which the rates of change of μ are significantly different.

As a side note, we bring up the fact that the two unstable contributions, associated with d_{cg} and d_b , are the same in magnitude if $J \equiv 1$. Indeed, using the definition of b , we observe that $\sum_{i=1}^m b^{i,i} := \nabla_\xi \cdot c$, where ∇_ξ denotes the Nabla operator (gradient) on unstable subspace. Thus, we can use Green’s first identity to

rewrite the latter term to

$$\begin{aligned} \int_M J d_b d\mu &= \int_M J \nabla_\xi \cdot c d\mu \\ &= - \int_M c \cdot \frac{\nabla_\xi(\rho J)}{\rho} d\mu \stackrel{J \equiv 1}{=} - \int_M c \cdot g d\mu, \end{aligned} \tag{28}$$

where ρ denotes the measure density conditioned on a local unstable manifold. It is now evident that the two ingredients of U , d_{cg} and d_b , involve both the array of $v-Q$ projections and a vector representing a local relative measure change. The only difference between them is that, in the latter term, the measure change is weighted by the value of J . If J is not strongly oscillatory nor has large gradients in phase space, then $\nabla_\xi(\rho J)/\rho$ behaves similarly to its non-weighted counterpart, g .

This analysis indicates that there are two possible ways of reducing the norm of U , either by manipulating c or g . According to the definition of c , reducing its norm would restrict our analysis only to a certain parameter. Note that c directly depends on χ , which represents the parametric perturbation of the trajectory. On the other hand, g contains information on the statistics of the unperturbed system. Therefore, neutralization of the effect of g might allow us to dramatically decrease $|U|$, regardless of the choice of a parameter with respect to which the linear response is computed. In the remainder of this section, the concept of “neutralization” will be explained in more detail.

Let us now consider a well-behaved objective function $J : M \rightarrow \mathbb{R}$, where M is an orientable compact manifold. Let the tangent bundle of M be expansive in all possible directions, which implies that all LEs are positive. Without loss of generality, we assume the volume integral of J over M is zero. Notice we can always add a constant number to J to ensure the zero mean condition, as the constant shift does not affect the linear response. Thus, J can be expressed in terms of the divergence of a vector field Z , i.e.,

$$J = \nabla_\xi \cdot Z. \tag{29}$$

After plugging Eq. 29 to the expression for U , we can apply Green’s first identity analogously to Eq. 28, which yields

$$U = - \int_M Z \cdot \frac{\nabla_\xi(\rho d)}{\rho} d\mu. \tag{30}$$

Note that Eq. 30 contains all combinations of mixed second derivatives of the SRB measure. To minimize the effect of the measure change, we want to eliminate

possibly as many components of g as possible, especially those corresponding to the least expansive directions. In an ideal scenario, we also want to neutralize the effect of those components of g that remain. This could be achieved by choosing a J that is *aligned* with q^1 , which means that the statistics of $\nabla_\xi \cdot Z = \sum_{i=1}^m \partial_{q^i} Z^i$ is dominated by its first term ($i = 1$), i.e.,

$$\|\partial_{q^1} Z^1\|_2 \gg \|\partial_{q^i} Z^i\|_2, \quad i = 2, \dots, m.$$

In this special case, we could approximate U by keeping only the first term of $\nabla \cdot Z$. For the truncated expression, we apply integration by parts, which yields

$$U \approx U^1 := \int_M \partial_{q^1} Z^1 d\mu = - \int_M Z^1 (d g^1 + \partial_{q^1} d) d\mu. \tag{31}$$

The first benefit of the alignment is that we automatically eliminate second differentiation with respect to the directions indicated by q^2, \dots, q^m that correspond to the largest slopes of μ . Therefore, the leading term of the unstable contribution is upperbounded as follows,

$$|U^1| \leq \|Z^1\|_2 \left(\|d g^1\|_2 + \|\partial_{q^1} d\|_2 \right). \tag{32}$$

The first term of the new inequality is proportional to $\|d g^1\|_2$. If $\|g^1\|_\infty \ll 1$, which is true if the measure is almost constant along q^1 , then $\|d g^1\|_2 \ll \|d\|_2$. This scenario is very likely in systems with a broad Lyapunov spectrum. In the second term of Ineq. 32, d is differentiated in the most expansive direction q^1 . It means that all components of the SRB density gradient weighted by c are differentiated once more. This time, however, we differentiate in the direction of the mildest descent/ascent of μ . One could visualize this process by considering the lateral boundary of a cylindrical solid. In this case, the tangent line computed along the solid’s height is always parallel to the solid and has zero slope. In any other direction, the slope is larger than zero. Differentiation of the nonzero slopes along the solid’s height effectively kills them all. We can apply this analogy to our case, in which we differentiate once more in the direction of the smallest slope. Therefore, the effect of the largest components of g corresponding to the least expansive directions could be neutralized, in which case $\|\partial_{q^1} d\|_2$ is expected to be negligible.

Through the above analysis, we conjecture that if J is aligned with the most expansive direction of the unstable manifold, as defined above, and the positive part of the Lyapunov spectrum is not clustered around

a certain value, it is possible to significantly reduce the magnitude of the unstable contribution. While the second condition is satisfied by many physical systems, the specific requirement for the objective function might be very restrictive. We now present a numerical example illustrating our argument.

In our investigation, we will focus on the following n -dimensional chaotic map $\varphi : [0, 2\pi]^n \rightarrow [0, 2\pi]^n$ defined as

$$x_{k+1}^i = 2x_k^i + s \sin(x_k^{i+1} - x_k^i) + t \sin(x_k^i) \bmod 2\pi, \quad i = 1, \dots, n, \tag{33}$$

where $n \in \mathbb{Z}^+, s \in \mathbb{R}, t \in \mathbb{R}$ and $x^{n+1} = x^1$. This is an extension of the one-dimensional *sawtooth map* [42], and therefore, we shall refer to φ defined by Eq. 33 as the *coupled sawtooth map*. The first term on the RHS introduces constant expansion that does not involve any parameters. Thus, if we set the coupling parameter to zero ($s = 0$), we obtain n independent maps with the same statistical behavior. If both the coupling and distorting terms are small, i.e., respectively, s and t are small, then all Lyapunov exponents are clustered around the value of $\log 2$, which means that the attractor is expansive in all directions. By increasing $|s|$, we strengthen the coupling between the neighboring degrees of freedom. For $n = 2$, the phase space gradient of the coupling term is parallel to the diagonal of the square manifold, $[0, 2\pi]^2$. Thus, the larger $|s|$, the stronger variations of the measure are expected along $[1, -1]^T$. In the case of a weak distortion, i.e., when $t \approx 0$, the SRB measure is expected to be approximately constant in the direction parallel to $[1, 1]^T$.

To verify these suppositions, we directly compute g for $n = 2$ at three different parameter sets: (1) $[s, t] = [0.05, 0]$ (weak coupling, no distortion), (2) $[s, t] = [-0.75, 0]$ (strong coupling, no distortion), (3) $[s, t] = [-0.75, 0.5]$ (strong coupling combined with distortion). For this purpose, we use a part of the full S3 algorithm to compute g along a trajectory (Lines 12–20 of Algorithm 2 in Appendix A) and plot both $|g^1|$ and $|g^2|$ on $[0, 2\pi]^2$. These results are illustrated in Fig. 6. In all three cases, the first component of g is statistically smaller in magnitude and features milder variations compared to the second one. They also confirm that the larger component of the relative measure change is approximately parallel to $[1, -1]^T$. Even in the presence of the distortion term (Case 3), the majority of white arrows, which indicate

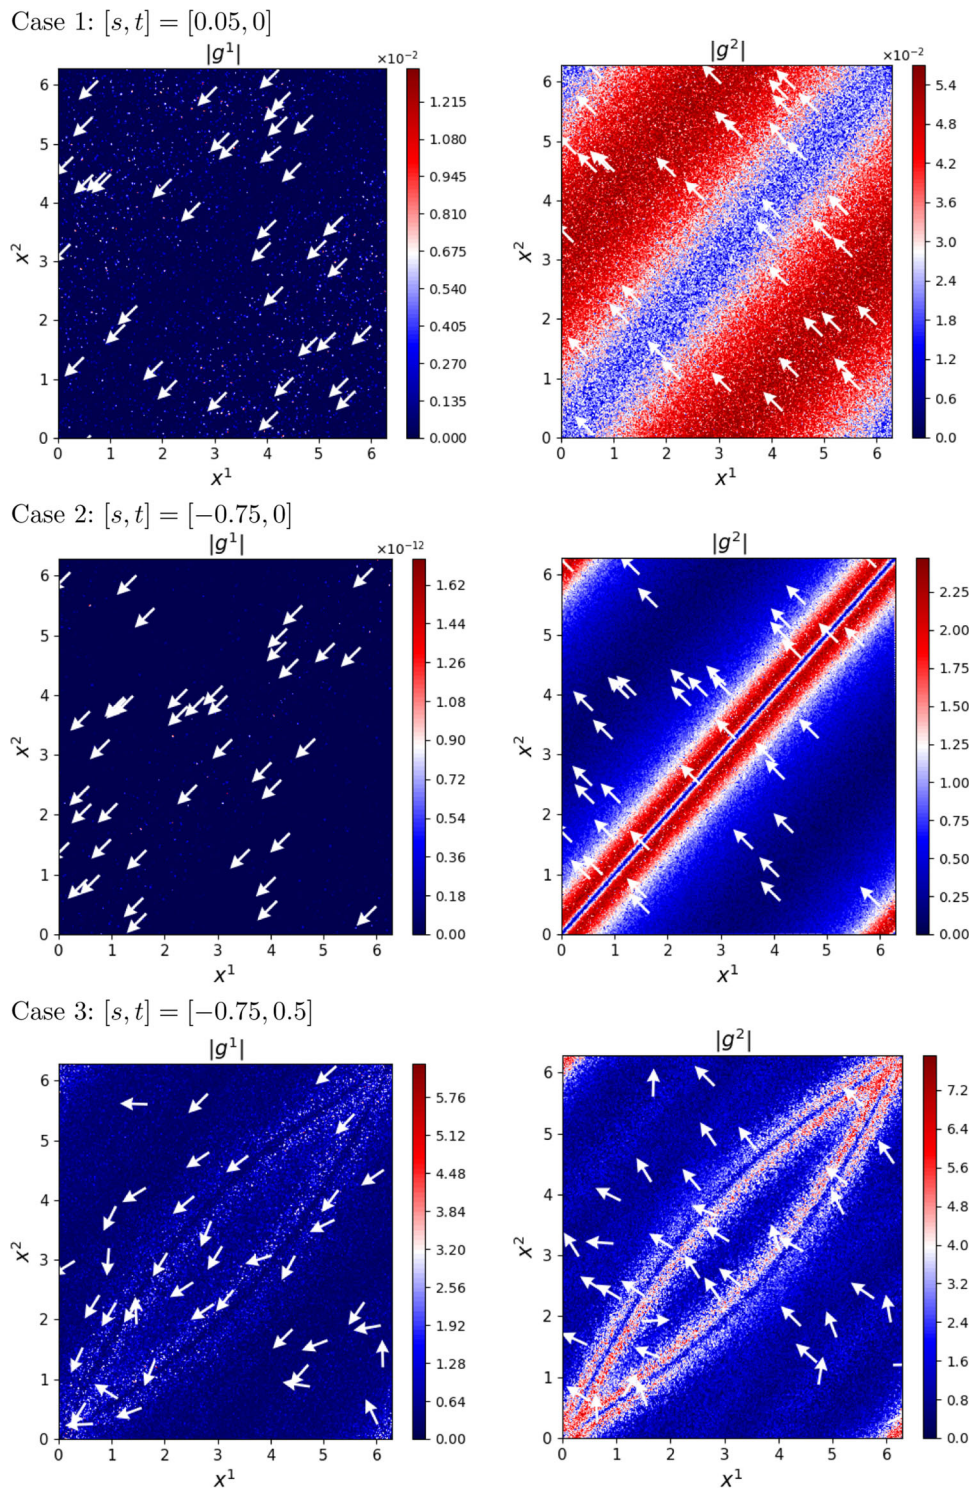


Fig. 6 Magnitude of both components of the SRB density gradient g of the two-dimensional coupled sawtooth map with two positive LEs. White arrows, respectively, represent q^1 and q^2 ,

which indicate local directions of differentiation. They are plotted every 5000 time steps. For each case, a trajectory of length $N = 3 \cdot 10^5$ was generated

local orthonormal directions q^1 and q^2 , tend to be oriented diagonal-wise. Notice that the larger coupling $|s|$, the larger rate of measure change in the least expansive direction represented by q^2 . If there is no distortion and coupling is significant (Case 2), then the first component of g is approximately zero everywhere in phase space. The largest measure gradients appear to be located around the $[1, 1]^T$ diagonal. Furthermore, if the coupling weakens, then the rates of expansion along q^1 and q^2 become similar. In Case 1, the distribution of g^1 has geometric features similar to its counterpart. This is consistent with our analysis, suggesting that both distributions are expected to have the same limits as $|s| \rightarrow 0$.

In Fig. 7, we plot the L^2 norms of selected components of g and corresponding Lyapunov exponents at different values of s and t . They were computed for the 2D ($n = 2$), 4D ($n = 4$), and 8D ($n = 8$) variants of the coupled sawtooth. In agreement with our conjecture, the norms of all components of g are equal and very small in the absence of the coupling term, i.e., when $s = 0$. We observe the norm ratio between g^1 and $g^m = g^n$ rapidly decreases as the coupling strengthens. This is also true between g^1 and other components corresponding to less expansive directions, as clearly indicated by the 4D and 8D examples. Figure 7 confirms the conjecture that the separation of Lyapunov exponents implies monotonic increase of the measure gradient norms as sorted from the most to the least expansive directions. Our results also indicate that if LEs are clustered around a single value, then the norm degradation is insignificant. Note that the converse is not necessarily true. Namely, there might be significant differences between particular components of g even if LEs are clustered, which is true for the 2D sawtooth map at $s \in [-1, 0]$. This usually happens when at least one of the components of g is no longer integrable with respect to μ [43]. We also acknowledge the fact that square-integrability of g with respect to μ is not required for the existence of the linear response, as we discussed in Sect. 2.3.

In light of the specific behavior of the SRB density gradient and our main conjecture presented above, we shall numerically investigate the impact of the objective function J on the statistics and their change with respect to parameters. The purpose of this experiment is to visualize long-time averages computed at different parameter values for the 2D coupled sawtooth map. A fundamental question we need to raise concerns the alignment requirement. How can we say that a chosen J is in fact aligned with q^1 ? Indeed, the two compo-

nents of the corresponding vector Z generally depend on both phase space coordinates. In the 2D setting, it is relatively straightforward to find a vector field Z that satisfies that requirement. If q^1 is approximately parallel to $[1, 1]^T$ and both components of Z depend on $x^1 + x^2$ only, i.e., $Z = Z(x^1 + x^2)$, the corresponding J is automatically aligned with q^1 . However, if $Z^1 = Z^1(x^1 + x^2)$ and $Z^2 = Z^2(x^1 - x^2)$, then their respective L^2 norms are expected to be similar. Finally, if $Z = Z(x^1 - x^2)$, then Z^2 becomes dominant giving more weight to the second component of g , which is in fact the least desired scenario.

Thus, we shall consider three wave-like objective functions that depend on $x^1 - x^2$, x^1 , and $x^1 + x^2$. These waves have zero gradients in the phase space directions parallel to $[1, 1]^T$, $[0, 1]^T$ and $[1, -1]^T$. They, respectively, represent functions that are weakly, moderately, and strongly aligned with the most expansive direction of the 2D hyperchaotic map. The statistics corresponding to these objective functions evaluated at a fine parametric grid are plotted in Fig. 8. We observe that the variation of statistics of $J = J(x^1 - x^2)$ is quite large in the regions that coincide with the parametric regime of a large measure change. Within this parametric subset, the value of the second LE evidently decreases and approaches the value of zero. Indeed, the largest sensitivity of the system is observed as s increases from $s \approx 0.35$ to $s \approx 0.5$ for all $t \in [-0.5, 0.5]$. Thus, for this parametric regime, the maximum value of $|d\langle J \rangle / ds|$ is $\mathcal{O}(1)$. In the moderate case, variations of $\langle J \rangle$ are significantly smaller compared to the previous example. However, we still observe non-negligible sensitivities of order $\mathcal{O}(10^{-1})$ if $s < -0.75$ and $|t| > 0$. The third plot of Fig. 8 shows the statistics of a function that is aligned with the most expansive direction, i.e., it depends on $x^1 + x^2$. The computed long-time averages now oscillate between two values that are $\mathcal{O}(10^{-3})$ apart, across the entire parametric space. These oscillations are distributed uniformly, even around the regions of large measure gradients and distortions. In this case, $\langle J \rangle$ is approximately independent of both parameters, which implies negligible linear response.

The major conclusion that follows from the above analysis and numerical examples is that the unstable part of the linear response might be negligible for a particular class of objective functions J . This is true for any system's parameter with respect to which the sensitivity is computed. We observed that a scattered distribution of the positive part of the LE spectrum leads

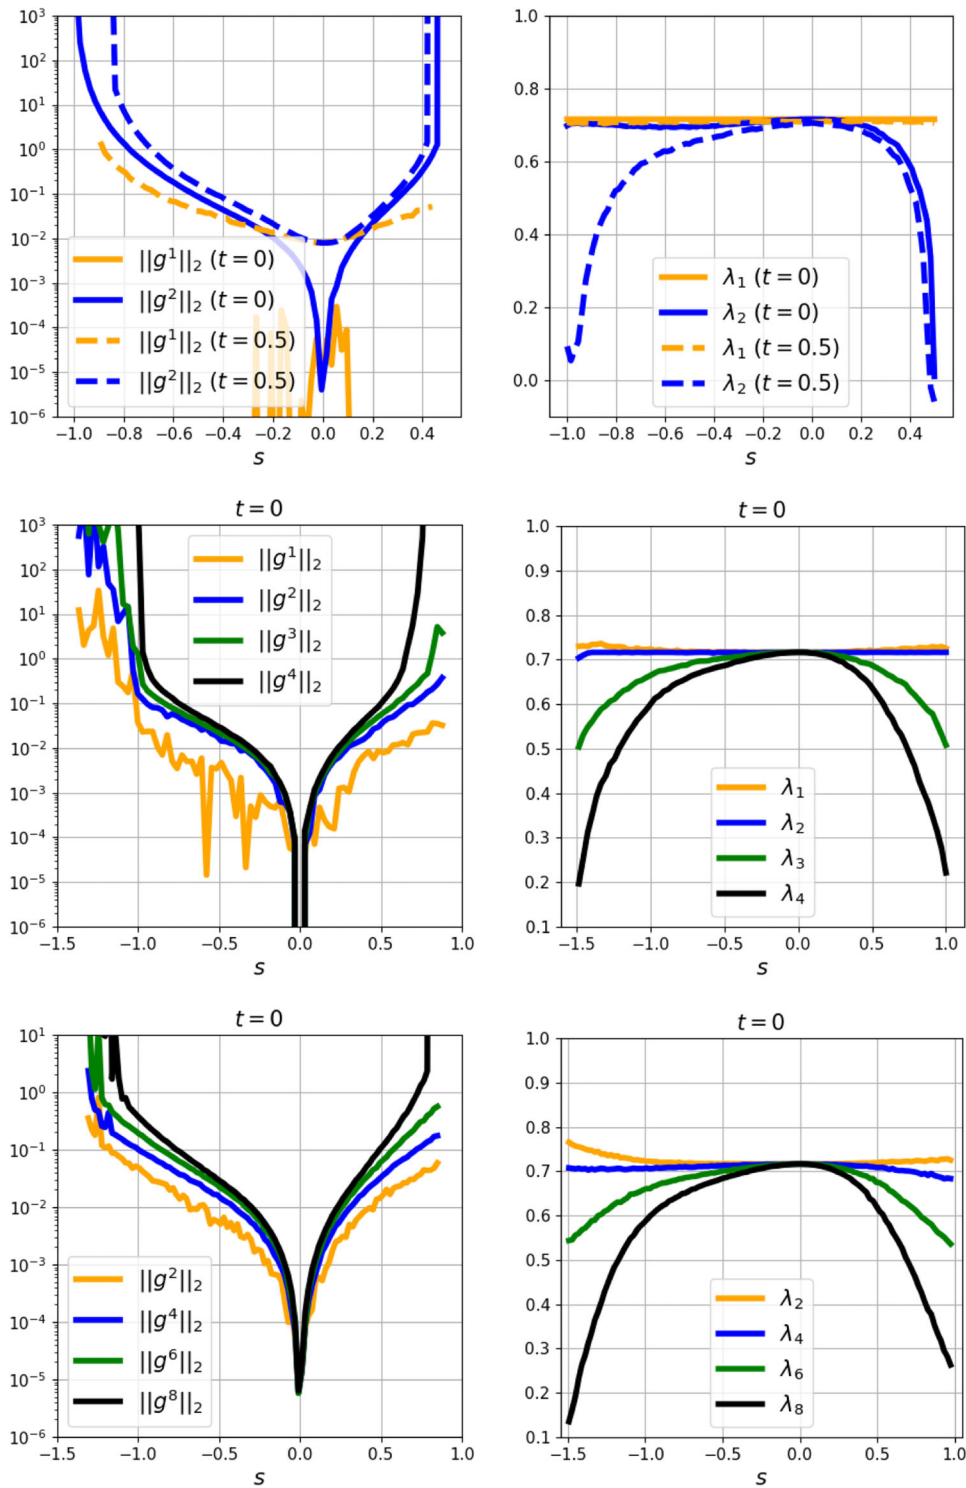


Fig. 7 L^2 norm of the SRB density gradient and Lyapunov exponents of the 2D ($n = 2$; top row), 4D ($n = 4$; middle row), and 8D ($n = 8$; bottom row) variant of the coupled sawtooth map.

All quantities were computed on a uniform grid of 100 values of the coupling parameter s . For each parameter, a trajectory of length $N = 3 \cdot 10^4$ was generated




Fig. 8 Long-time averages of the wave-like objective function $J = \exp(\sin(z)) \sin(z)$, where $z = x^1 - x^2$ (upper plot), $z = x^1$ (middle plot) and $z = x^1 + x^2$ (lower plot). The time averages were computed for a uniform parametric grid consisting of 225 and 100 points along s and t , respectively. For each set of parameters, a trajectory of length $N = 5 \cdot 10^6$ was generated. The dashed lines represent isolines corresponding to two different values of the second (i.e., smaller) LE: 0.5 (dark blue) and 0 (violet)

to the norm increase of consecutive components of the SRB measure gradient, represented by g . This usually causes significant variations of the statistics in the parametric space and, simultaneously, enables finding the optimal alignment of J . In this section, we demonstrated that the elimination/neutralization of the largest components of the SRB measure gradient might dramatically reduce the unstable contribution. This can be achieved by choosing a J that is aligned with the most expansive direction, which is reflected by the partial integration in Eq. 31. In high-dimensional systems, we expect substantial reductions of the unstable contribution as long as J is aligned with any subspace spanned by the most expansive directions. Note also that our argument applies only to systems with at least two positive LEs. If $m = 1$, there is only one expansive direction, which means there are no degrees of freedom for choosing an appropriate J .

How can these results and analysis be used in the context of practicable high-dimensional systems? In a standard engineering design process, the quantity of interest is a well-defined function with a concrete physical meaning, e.g., temperature, kinetic energy, drag force, that is generally not aligned with some abstract subspace of the chaotic attractor. In the following section, we argue that the specific condition imposed on J is not an obstacle for a vast family of dynamical systems encountered in many fields such as climate science and turbulence theory. We show that the stable part alone can approximate the total linear response sufficiently well.

4 Sensitivity analysis of higher-dimensional flows with statistical homogeneity

We presented an argument supporting the concept of small unstable contributions. This promising observation may lead to a significant simplification of the S3 algorithm for the linear response. As described in Sect. 3, the major requirement for the leading unstable term U to be small is a concrete alignment of the objective function J . In an ideal setting, the slope (variation)

of J in the least expansive directions should be relatively low compared to the most expansive one represented by q^1 . This requirement seems to be very restrictive given complicated dynamical behavior of general high-dimensional chaos. In the simple example introduced in Sect. 3, the most expansive direction was predictable, thanks to which one could easily choose a suitable J . In this section, we will focus on a common feature of a vast group of spatially extended chaotic systems: statistical homogeneity in space. Relying on this property, we argue that the system’s dimension n increases the probability of the desired alignment, regardless of the physical meaning and form of J .

Statistical homogeneity in the physical space implies that the long-time behavior of all system coordinates is approximately the same. For such systems, the objective function is usually defined in terms of the spatial average of a physical quantity. For 1D-in-space continuous systems bounded by $a \in \mathbb{R}$ and $b \in \mathbb{R}$, $b > a$, for example, J is usually expressed as follows,

$$\begin{aligned}
 J &= \frac{1}{b-a} \int_a^b \tilde{J}(x) dx \approx \frac{1}{n} \sum_{i=1}^n \tilde{J}(x^i) \\
 &:= \frac{1}{n} \sum_{i=1}^n \tilde{J}^i \tag{34}
 \end{aligned}$$

where $\tilde{J} : \mathbb{R} \rightarrow \mathbb{R}$ is a function with a concrete physical meaning. In the case of the Navier–Stokes model, \tilde{J} is linear if the velocity is the quantity of interest. For energy-like quantities, such as the kinetic energy, \tilde{J} could be a quadratic function. Note that if the property of statistical homogeneity holds, then

$$\langle J \rangle = \langle \tilde{J}^1 \rangle = \langle \tilde{J}^2 \rangle = \dots = \langle \tilde{J}^n \rangle,$$

where $\langle \cdot \rangle$ denotes the long-time average. This implies that for any time-dependent weight vector $w(t) \in \mathbb{W}$, where

$$\mathbb{W} = \left\{ w \in \mathbb{R}^n \mid \sum_{i=1}^n w^i(t) = 1 \ \forall t \geq 0 \right\},$$

the following is true

$$\begin{aligned}
 \langle J_w \rangle &:= \left\langle \sum_{i=1}^n w^i \tilde{J}^i \right\rangle = \sum_{i=1}^n \langle w^i \tilde{J}^i \rangle \\
 &\stackrel{\text{indep.}}{=} \langle \tilde{J}^1 \rangle \left\langle \sum_{i=1}^n w^i \right\rangle = \langle J \rangle.
 \end{aligned} \tag{35}$$

Equation 35 assumes \tilde{J}^i and its corresponding weight are statistically independent. Therefore, the original

objective function J can be replaced by any member from the class of spatially weighted functions without affecting the long-time behavior. This critical observation implies that for any smooth J , the feasible space of J_w increases with the system’s dimension n . It means that for a large n , there might be a lot of candidates well-aligned with q^1 . Note that w should primarily depend on q^1 , i.e., an inherent topological property of the tangent space, which justifies the assumption of statistical independence of w and a single phase space coordinate and, consequently, independence of \tilde{J}^i and w in the limit $n \rightarrow \infty$.

We highlight yet another common property of larger physical systems. As reported by several publications (see [32] and references therein), one can distinguish spatially localized structures of the expansive part of the covariant Lyapunov basis. For example, in a 3D turbulent flow past a cylinder studied in [28], the most expansive directions tend to be localized in the areas of primary instability. These include the boundary layers and near-wake regions. In far-wake regions and in the free stream, the most expansive (leading) covariant Lyapunov vector (CLV) was reported to be inactive, i.e., approximately zero. Moving away from the regions of primary instability, less expansive and contracting CLVs tend to be dominant. However, as pointed out in [32], in homogeneous systems with periodic boundary conditions, the clustered activity regions of the leading CLV may move across the entire physical domain. In their analysis of Rayleigh–Bénard convection [48], the authors notice that, for the most expansive CLVs, the energy spectral density is concentrated around a specific wave number, which turns out to be approximately the same as the one of the primal solution. The same work demonstrates that the energy spectrum density gradually becomes uniform as the CLV index increases. Based on the rich numerical evidence, we expect that any time instance q^1 is expected to involve local activity patterns that are restricted to a sub-region or wobble around the entire domain. Recall that q^1 and the leading CLV are the same up to a multiplicative prefactor. This is no longer true for q^i , $i = 2, \dots, n$, due to the orthonormalization procedure.

Given these specific properties of higher-dimensional chaos, the problem of alignment of J and q^1 could be easily circumvented. Notice that we have freedom in choosing time-dependent weights, which can potentially favor only those coordinates that correspond to the regions of “activity” of q^1 . As these “activity” clus-

ters move around in time, the corresponding weights can be adjusted accordingly keeping the remaining components of w close to zero. If $\tilde{J}^i = x^i$, then the optimal choice of weights is strictly determined by the components of q^1 . For higher-order polynomial objective functions, the relative values of state components would also affect the corresponding weights. Their individual contributions, however, are negligible if n is large. A high density of spatial coordinates facilitates search of the optimal set of weights favoring the active components of J in the right proportion, regardless of the form of \tilde{J}^i . For a dynamical system with arbitrary statistical behavior and complex tangent topology, it is generally difficult to analytically estimate how large n should be to ensure the satisfactory alignment of J_w leading to the neutralization of the unstable term. Therefore, in this section, we resort to numerical studies of systems with statistical homogeneity to guarantee that Eq. 35 holds.

Before we discuss the numerical results, we first focus on algorithmic consequences of neglecting the effect of the SRB measure change. Indeed, a complete omission of the unstable part in the computation of linear response dramatically simplifies the space-split algorithm. That term, obtained through partial integration, requires computing the SRB density gradient and derivatives of projections of tangent solutions onto the unstable-center subspace. These two ingredients require solving $\mathcal{O}(m^2)$ second-order tangent equations, which is by far the most expensive section of Algorithm 2. Assuming n is large, further simplifications can be introduced. Note that the neutral contribution involves an infinite series of k -time correlations of c^0 and $DJ \cdot f$ with the leading term

$$C = \int_M c^0 DJ \cdot f \, d\mu := \int_M (c^0 |f|) DJ \cdot q_f \, d\mu, \tag{36}$$

where c^0 is the projection of a center-stable component of the tangent solution onto the center subspace normalized by the length of f as derived in Eq. 11. Notice that the form of C is in fact identical to its unstable counterpart in its original form. Therefore, if our conjecture of small unstable contributions applies, then C is also small and can be neglected in the linear response algorithm. Indeed, the L^2 norms of $DJ \cdot q_f$, $DJ \cdot q^u$ are expected to be similar, where q^u is some unstable direction, unless the positive Lyapunov spectrum is clearly bounded away from zero. Recall also that the

projection coefficients c^i , $i = 0, 1, \dots, n$ represent dot products of a component of v and their corresponding tangent vectors. The direction of parametric deformation is generally independent of Lyapunov vectors. We later demonstrate that these coefficients become similar in value as $n \rightarrow \infty$. Based on this analysis, we conclude that if our conjecture of a small U holds, then the computation of C could also be neglected.

Exclusion of both unstable and neutral terms from the full S3 algorithm leaves us with the stable term alone. The remaining part requires computing the regularized tangent solution through step-by-step orthogonal projection of the unstable-center component. Since f is generally not orthogonal to the column space of Q , the original stabilizing procedure involves an assembly and inversion of the Schur complement S . We have directly used f because it is always given at no cost and it allows for a straightforward derivation of a computable formula for the neutral part of the linear response. However, since we neglect that part as well, the process of regularizing the tangent solution can be simplified even further. Instead of using f and then orthogonalizing the (Q, f) tuple, we can solve one more first-order tangent equation and perform QR factorization of the extended tangent solution matrix. Thanks to this modification, we recursively generate the orthogonal basis of the unstable-center subspace and compute projections of v onto that basis, which is equivalent to the original algorithm. This can be achieved by executing Lines 9–10 of Algorithm 2 by changing m to m_{ext} , where m_{ext} should ideally be equal to $m + 1$. In practice, however, setting $m_{ext} = m + 1$ may lead to instabilities due to the potentially non-hyperbolic behavior of the system. Moreover, if n is large, we rarely know the exact value of m . If our aforementioned conjecture of a small C is valid for large systems, then we could project out a few additional components of the tangent space from v . Therefore, as long as m_{ext} is close to $m + 1$, the penalty of these extra projections, in the context of sensitivity approximation, is expected to decrease as $n \rightarrow \infty$. The only practical consequence is that a few extra tangent equations will have to be solved, which barely influences the overall cost of the reduced algorithm assuming $m_{ext} - m \ll m$. Algorithm 1 summarizes all steps required to approximate the sensitivity. This procedure was obtained by eliminating the unstable and neutral contributions from the full S3 algorithm. By-products of the S3 algorithm are Lyapunov exponents, included in the le array, which

we compute to supplement our discussion. Benettin in [4] originally proposed this approach for approximating LEs.

4.1 Lorenz 96

In light of the above conclusions, we shall consider the Lorenz 96 model, which was proposed by E. Lorenz in [25], to study spatiotemporal dynamics of the atmosphere. Mathematically, this is an n -dimensional chaotic flow defined as follows,

$$\begin{aligned} \frac{dx^i}{dt} &= (x^{i+1} - x^{i-2})x^{i-1} - x^i + F, \quad i = 1, \dots, n, \\ x^{i+n} &= x^i, \end{aligned} \quad (37)$$

where the superscript indicates the component index, in compliance with our notation convention. Each degree of freedom x^i represents a value of a physical quantity, e.g., temperature or pressure, on a uniformly discretized parallel of the Earth. Analogously to semi-discretized PDEs describing advection, this system involves spatially coupled variables with a quadratic nonlinearity. Equation 37 involves two constant parameters: the number of sectors $n \geq 4$, each corresponding to a different meridian of the Earth, and imposed forcing $F \in \mathbb{R}^+$. If $F < 8/9$, then the solution quickly decays to the constant value of F , i.e., $x^i = F$, $i = 1, \dots, n$ for all $t > t^* \approx 0$ [17]. We solve Eq. 37 using the explicit fourth-order Runge–Kutta with $\Delta t = 0.005$. That ODE solver will be used throughout this section, unless stated otherwise. In Fig. 9, we plot the solutions for $n = 80$ and three different values of F . For $F = 3$, the periodic dynamics involves waves travelling to the west, i.e., in the direction of decreasing sector index i . The distortion that appears at the beginning of the simulation quickly decays leading to a predictable behavior. While some regularity is still maintained at $F = 6$, the alignment of waves seems random, which implies that some unstable modes might be activated. If we further increase F to the value of 9, the spatiotemporal structure of the solution clearly reflects chaotic behavior without any distinguishable patterns.

To obtain more insights into the dynamics of the Lorenz 96 model, we analyze its Lyapunov spectrum for the most common values of the system's parameters [45]. In Fig. 10, we illustrate a half of the Lyapunov spectrum for $F \in [0, 25]$ at $n = 10, 20, 40, 80$. For any

n and $F < 0.9$, all LEs are negative, which means that, for any random initial condition, the solution exponentially decays to a constant value. Within the interval $F \in [0.9, 4.5]$, the dynamics is no longer stationary, but still non-chaotic, because $\lambda_1 = 0$. We observe the presence of at least one positive LE if $F > 4.5$. In the chaotic regime, the dimension of the expansive manifold gradually increases with F to about $m = n/2$ at $F = 25$. Notice also that the higher F , the smaller the angle between the lines representing $\lambda_i(F)$, $i = 1, 2, \dots$ and the x-axis. Indeed, the authors of [17] computed a curve fit for $\lambda_1^{-1}(F)$ at $N = 35$, whose close-form formula is the following: $\lambda_1^{-1}(F) = 0.158 + 123.8 F^{-2.6}$. Consequently, given the self-similar behavior of the plotted spectrum, all LEs seemingly converge to fixed values as the forcing F increases.

We shall consider the spatially averaged kinetic energy of the system as the objective function J , which can be expressed using Eq. 34 with $\tilde{J}^i = (x^i)^2$. The long-time averages $\langle J \rangle$ for $F \in [0, 25]$ at $n = 10, 20, 40, 80$ are plotted in Fig. 11. We observe that all four curves $\langle J \rangle(F)$ collapse into a single curve due to spatial averaging. The only misalignment occurs at the non-chaotic/chaotic transition region close to $F = 5$. Thus, in the extensive chaos regime of Lorenz 96, the spatially averaged statistics is generally independent on n , which was previously observed in [17]. We shall restrict our attention to that regime, i.e., when $F \geq 5$, and compute sensitivities with respect to F using our reduced S3 algorithm. The slope of $\langle J \rangle(F)$ seems to be constant and is approximately 2 for $F \in [5, 25]$. We will use a higher-order interpolation of the statistics curve and differentiate it using the central finite-difference scheme. This estimate will serve as a reference solution to evaluate the performance of Algorithm 4.


Figure 12 illustrates approximations of the linear response obtained with Algorithm 1. In particular, we used our reduced algorithm to approximate $d\langle J \rangle/dF$ for $F \in [5, 25]$. For $m_{ext} = m + 1$, the algorithm generates satisfactory approximations for $F \geq 6$. However, the standard deviation is quite large, and it very often exceeds the value of one across the entire parametric domain. These statistical fluctuations are eliminated by increasing m_{ext} . Indeed, the $m_{ext} = m + 2$ case has dramatically smaller sigmas everywhere. This result indicates that if m_{ext} is too small, the regularized tangent solution may still have rapidly growing com-

Algorithm 1 Reduced space-split sensitivity algorithm for higher-dimensional chaotic flows

Input: N, K, T, n, m_{ext}
Output: $d\langle J \rangle / ds \approx s / N, m_{ext}$ largest LEs $:= le / N$

- 1: Randomly generate: x_0, v_0, Q_0 such that $\text{size}(x_0) = \text{size}(v_0) = (n, 1), \text{size}(Q_0) = (n, m_{ext})$
- 2: Set $s = 0$ and $le = \text{zeros}(m_{ext})$
- 3: **for** $k = 0, \dots, N - 1$ **do**
- 4: **if** $k \geq T$ **then**
- 5: $s := s + DJ_k \cdot v_k$
- 6: $le := le + \text{diag}(\log(\text{abs}(R_k)))$
- 7: **end if**
- 8: $P_{k+1} = D\varphi_k Q_k$
- 9: QR-factorize P_{k+1} : $Q_{k+1} R_{k+1} = P_{k+1}$
- 10: $r_{k+1} = D\varphi_k v_k + \chi_{k+1}$
- 11: $c_{k+1} = Q_{k+1}^T r_{k+1}$
- 12: $v_{k+1} = r_{k+1} - Q_{k+1} c_{k+1}$
- 13: Advance the iteration: $x_{k+1} = \varphi(x_k)$
- 14: **end for**

Fig. 9 Solutions to the Lorenz 96 system (Eq. 37) for $n = 80$ stacked horizontally



ponents in some parts of the attractor leading to large variances. The smooth behavior of the linear response in the $m_{ext} = m + 2$ case suggests that these fluctuations are not caused by the ergodic-averaging error. As expected, there is always an extra penalty for increasing m_{ext} . However, the higher n , the smaller price must be paid for extra stabilizing projections. This observation is consistent with our conjecture, suggesting that the relative contribution of a single component of v decreases as n gets larger.

Figure 12 reveals two other critical features of the reduced algorithm. First, if n is sufficiently large, then the obtained sensitivity approximation might be very accurate, i.e., the relative error is no larger than a few

percent. This result confirms our major conjecture of negligible unstable (and neutral) contributions to the total linear response. For Lorenz 96, the impact of the SRB measure change is apparently insignificant. The only exception is the region around $F = 5$. Indeed, the error is large in this parametric regime, regardless of the value of m_{ext} and system’s dimension n . Although the property of spatial homogeneity is unaffected and some unstable modes are still active, we observe the sensitivity approximation clearly deviates from the reference solution. Note that this parametric region coincides with the rapid decrease of positive LEs. Many of them are still positive, but they are clustered. Our discussion in Sect. 3 suggests that in this case there might be no




Fig. 10 Larger half of the Lyapunov spectrum of Eq. 37. LEs were computed at 240 distinct values of F distributed uniformly between 0 and 25. For each value of F , we run 10 independent simulations over 5000 time units. The barely visible shaded area represents the 2-sigma range (95% confidence) of the 10-element data set at each value of F




Fig. 11 Long-time means of spatially averaged kinetic energies of the Lorenz 96 system. The statistics were computed on a uniform grid of 240 values of $F \in [0, 25]$. For each value of F , the objective function was time-averaged over $5 \cdot 10^6$ time units

gain due to the alignment of J and q^1 . All components of g are expected to have similar distributions across the phase space. Therefore, even if J and q^1 are aligned, the unstable contribution could be significant in this case.

For completeness, in Fig. 13, we also plot the L^2 norms of the projection scalars $c^i, i = 1, \dots, m_{ext} = m + 2$. This result confirms that all scalars contribute almost equally to the linear response, suggesting that their relative significance is degraded as n increases. These results also indicate that if n is small, the scalars corresponding to the lowest indices tend to be statistically larger compared to their counterparts. In other words, the Lorenz 96 system with few degrees of freedom tends to favor the contributions of $\|c^i\|_2$ corresponding to the most expansive directions.

4.2 Kuramoto–Sivashinsky

Finally, we shall consider the Kuramoto–Sivashinsky (KS) equation, one of the simplest partial differential equations modeling chaos. Similarly to Lorenz 96, KS is a spatiotemporal description of complex dynamics driven by instabilities far from an equilibrium. This equation was proposed decades ago to model wave propagation in reaction-diffusion systems [21] and hydrodynamic instabilities of laminar flames [38]. A number of other applications of the KS equation can be found in the literature. In this work, we analyze a modified version of KS, which includes an extra advection term proportional to a constant scalar $c \in \mathbb{R}$. The modified equation, which was previously studied in [6],




Fig. 12 Linear response approximations of the Lorenz 96 model with respect to F computed using Algorithm 1. The top plot illustrates approximated sensitivities for $m_{ext} = m + 1$, the middle plot for $m_{ext} = m + 2$, while the bottom plot depicts the mean relative error of the $m_{ext} = m + 2$ case computed with respect to the reference finite-difference solution (respective colors indicate n). Sensitivities were computed on a uniform 240-point grid between $F = 5$ and $F = 25$. For each value of F , we run 10 independent ergodic-averaging simulations over $N\Delta t = 5000$ time units. Vertical lines represent sigma intervals, while the bullets indicate the corresponding averages. Lack of a bullet (in the upper plot) means the standard deviation is larger than 1. The solid orange line is a finite difference approximation of the 11-th degree polynomial fit of $\langle J \rangle$




Fig. 13 L^2 norms of c^i , $i = 1, \dots, m_{ext} = m + 2$, which were computed as by-products of Algorithm 1. All simulation parameters are the same as those reported in the caption of Fig. 12

has the following form,

$$\begin{aligned} \frac{\partial u}{\partial t} &= -(u + c) \frac{\partial u}{\partial x} - \frac{\partial^2 u}{\partial x^2} - \frac{\partial^4 u}{\partial x^4}, \\ u(0, t) &= u(L, t) = 0, \\ \frac{\partial u}{\partial x}(0, t) &= \frac{\partial u}{\partial x}(L, t) = 0, \end{aligned} \tag{38}$$

where $x \in [0, L]$, $L = 128$, $t \geq 0$, $u(x, t) \in \mathbb{R}$. We discretize this system in space using the finite difference method with second-order accuracy. The grid is uniform and involves 513 nodes, which gives us a constant spacing $\Delta x = 128/(513 - 1) = 0.25$. A combination of center and one-sided schemes is applied to approximate all spatial derivatives as suggested in [6]. The number of ODEs, i.e., the system's dimension, is reduced to $n = 511$ by incorporating all boundary conditions using the ghost node technique. While this is a stiff system, we apply the fully-explicit

fourth-order Runge–Kutta scheme with a small time step $\Delta t = 0.0006$. In Appendix B, we discuss how the linear response algorithm could be integrated with implicit schemes.

Figure 14 illustrates solutions to the KS equation, $u(x, t)$, for different values of c . In the spatiotemporal space, $u(x, t)$ involves a collection of irregular branches that switch between positive and negative values. The sign of c determines the inclination of these branches. If c is positive, they tend to move in the positive direction of x and vice versa. By increasing the magnitude of c , the advection term starts to dominate pushing the lightly turbulent region out of the domain. Indeed, for $c = 2$, we observe that $u(x, t)$ quickly becomes steady suggesting that all unstable modes are killed due to the strong advection. Regardless of the value of c , one can distinguish a transitional period at the beginning of each simulation during which the spatiotemporal branches develop their shapes. At $c = 1.4$, the spatial sub-region $x < 20$ is dominated by the convection, which results in an almost stable behavior of $u(x, t)$ in that part of the domain. This leads to violation of statistical homogeneity along x .

Figure 15 depicts the 18 largest Lyapunov exponents of the KS equations for $c \in [-1, 2]$. The LE spectrum is independent of c as long as $-1 \leq c \leq 1.3$. At $1.3 \leq c \leq 1.7$, we observe a rapid decrease of all positive LEs. This coincides with the increasing strength of the advection term. Intuitively, the dominating advection term gradually kills the unstable modes, which consequently leads to a more predictable behavior of $u(x, t)$. The KS system is clearly non-chaotic if $c > 1.7$, which is reflected by the stable behavior of $u(x, t)$ at $c = 2$ illustrated in Fig. 14.

We also acknowledge similarities in the behavior of LE spectra corresponding to the Lorenz 96 and KS system. In the former, we observed an analogous collapse of the values of positive LEs around the laminar-to-turbulence transition close to $F = 5$. Another analogy is the parametric independence of the LE spectrum at large values of F . Note, however, that the ratio m/n may reach the value of $1/2$ in the case of Lorenz 96, which is significantly larger compared to this case.

Selected Lyapunov vectors are plotted for $t \in [0, 1200]$ in Fig. 16. As expected, the leading Lyapunov vector q^1 consists of relatively large structures with local support. The region of activity of q^1 , which corresponds to non-small components, is limited to a thin sub-region, which moves around the entire x -space. It

periodically bounces back and forth between the two walls. We observe that the structural behavior of q^i visibly changes as i increases. The support of q^{20} is rather global with occasional small inactivity regions. The same is true for q^{40} , which also features much finer structures compared to the previous two. The q^{60} vector, on the other hand, seems to be periodic and highly oscillatory in x , and almost constant (stationary) in t across the entire spatiotemporal domain. The tangent vectors corresponding to moderate indices are placed in the bottom row of Fig. 16. They consist of finer structures compared to the ones of q^1 and have occasional small inactivity regions throughout the entire domain. All vectors in the bottom row are visibly similar except when t is small. Recall that all Lyapunov vectors q^i were obtained in an iterative procedure involving a set of forward tangents that is initiated at a random initial condition. We observe that this iteration persistently requires at least 50 time units for a convergence run-up.

We also highlight the fact that, due to the recursive orthonormalization procedure, several physical features are lost. While the orthogonal Lyapunov vectors are sufficient to determine a basis of unstable or center-unstable subspaces required for our linear response algorithms [4, 40], they cannot be directly used to compute the individual contractive or center directions of the tangent space, nor can they be used to approximate the angles between different tangent subbundles. Hence, more information is required to study the hyperbolicity of a system [20, 22, 44].

Given these preliminary results, we apply Algorithm 1 to compute linear response with respect to the parameter c . This time we shall consider three different spatially averaged objective functions: linear, quadratic and cubic, i.e., $\tilde{J}^i = u^p$, $p = 1, 2, 3$, respectively. The corresponding long-time averages are plotted in Fig. 17 at $c \in [-1, 2]$. We observe that, in all of these cases, the mean curve can be divided into three smooth sections connected at $c \approx 1.25$ and $c \approx 1.7$. The shape of the left part resembles a polynomial function of the same order as the objective function itself. The middle one resembles the tangent function, while the right-hand side piece is constant in all three cases. These three pieces coincide with three different behavior types of $u(x, t)$ that we observed in Fig. 14: turbulent ($c \leq 1.25$), transitional ($1.25 \leq c \leq 1.7$), and advection-dominated ($c \geq 1.7$) regime.

We apply our reduced linear response algorithm (Algorithm 1) to approximate sensitivities for these

Fig. 14 Solutions to the KS equation (Eq. 38) for different advection intensities





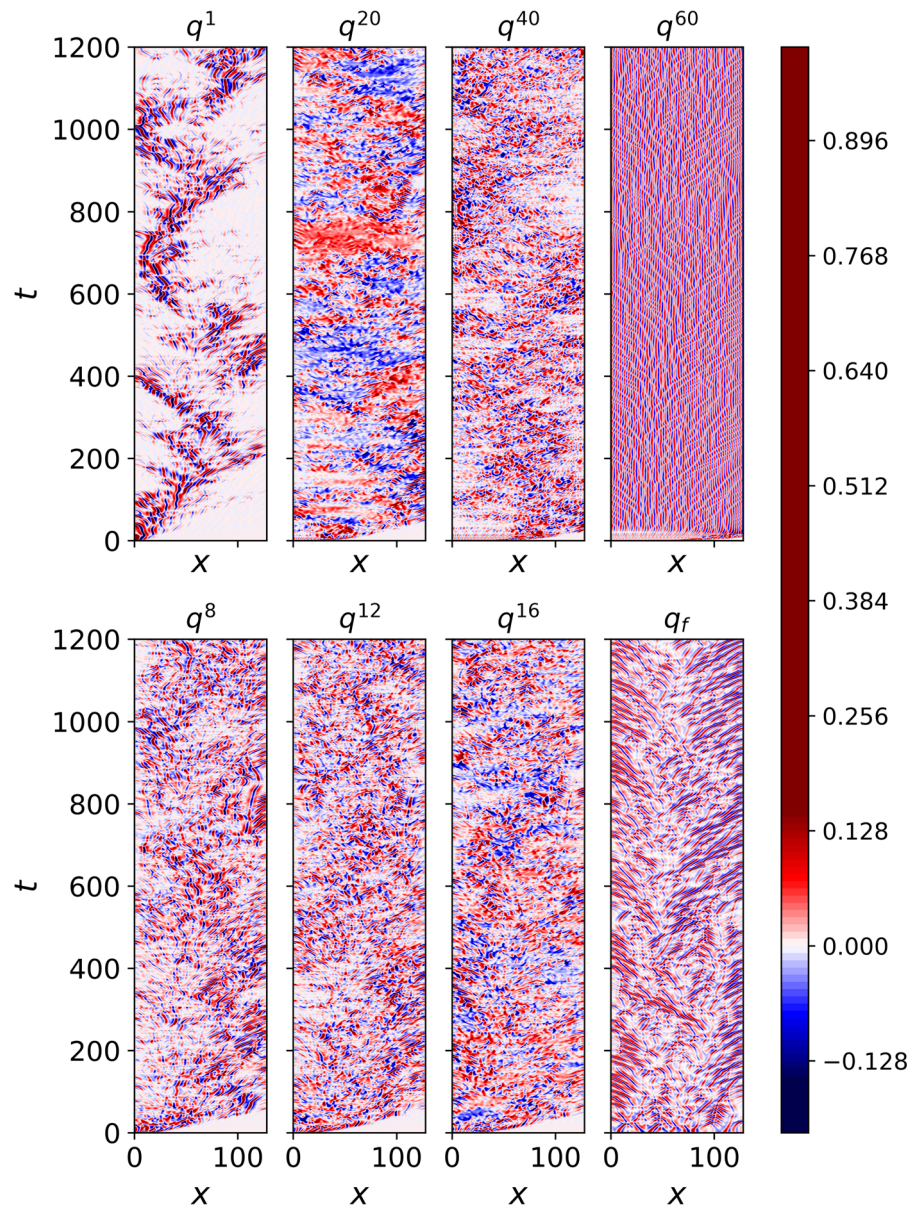
Fig. 15 18 largest Lyapunov exponents of the KS equation. The spectrum was computed at the uniform grid between $c = -1$ and $c = 2$. For each value of c , 10 independent simulations were run. The sought-after quantities were obtained through ergodic-averaging over 12,000 time units per simulation. The solid lines represent the mean values obtained in 10 simulations, while the shaded area represents the 2-sigma range



three objective functions. Analogously to the previous plots, we compare our approximations against the finite-difference reference solutions. Figure 18 illustrates the linear response results for different values of m_{ext} . One can easily observe a lot of similarities between these results and the ones generated for Lorenz 96. First of all, if $m_{ext} = m + 1$, the mean solution is quite close to the reference line, but the variance is likely to be large. The variance is significantly reduced by increasing m_{ext} and, in most cases, the new

mean approximations are still very accurate. Indeed, the accuracy can be within the reference line width in the turbulent and stable regimes. Huge disparities occur in the transitional regime, i.e., at $c \in [1.25, 1.7]$. Similarly to the Lorenz 96 case, this region corresponds to the sudden decrease of positive LEs. The approximation errors here are generally smaller compared to those computed for the Lorenz 96 system. Recall that, in Fig. 12, we observed that the approximation error decreases as $n \rightarrow \infty$. Indeed, the dimension of the

Fig. 16 Orthonormal Lyapunov vectors q^i of the KS system (Eq. 38) without the extra advection term ($c = 0$). The vector q_f represents the normalized time derivative of $u(x, t)$. The colorbar has linearly been re-scaled between -0.15 and 0.15 keeping the same color for all values from beyond this interval



discretized KS system is an order of magnitude larger than that of Lorenz 96.

Our numerical results presented in this section indicate that the linear response of a higher-dimensional system can be accurately approximated by the reduced S3 method. That algorithm, which was obtained by eliminating the unstable and neutral contributions, solves a regularized tangent equation by projecting out all expansive and, if necessary, a few other tangent components. This process can be in fact formulated as an optimization problem in which we minimize the L^2

norm of the sum of the standard tangent solution and a linear combination of expansive orthogonal Lyapunov vectors. A similar concept was previously utilized in a variant of shadowing methods known as NILSS [31], which relies on covariant Lyapunov vectors. While there are some algorithmic differences between the reduced S3 and NILSS, this work also sheds light on the reliability of relatively simple methods using some form of a regularized tangent equation.

We also note that there is potential in applying the reduced version of the linear response algorithm to the




Fig. 17 Long-time averages $\langle J \rangle$ computed on a uniform 240-point grid of $c \in [-1, 2]$. The operator $\langle \cdot \rangle_x$ indicates the spatial average. For each value of c , we run an ergodic-averaging simulation over 600,000 time units

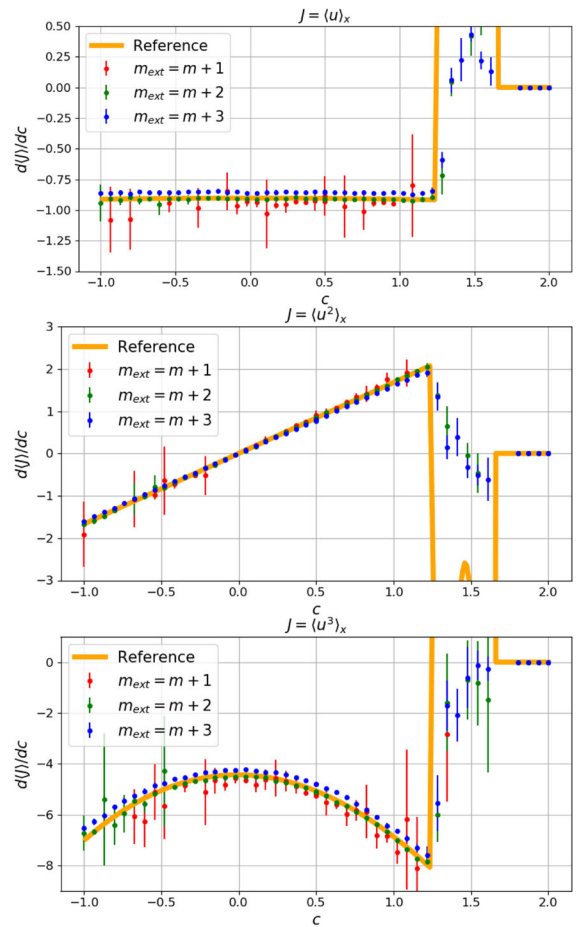


Fig. 18 Linear response computed for the same objective functions as those presented in Fig. 17 using Algorithm 1. For each value of c , we run 10 independent simulations over 3,000 time units each. Bullets and vertical lines represent the mean and standard deviation, respectively. The results with a large standard deviation were removed from the plot. The reference line was computed through central finite-differencing of polynomial fits

broad family of time-delayed dynamical systems. The spatiotemporal structure of the laser dynamics with delayed feedback presented in [2, 15] clearly features a statistically homogeneous behavior. The user would need to represent such a system using an appropriate diffeomorphic map $\varphi : M \rightarrow M$ and compute relevant

phase-space and parametric derivatives, following the recipe described in this paper. For systems with delay τ and constant time step Δt , one can consider introducing approximately $\tau/\Delta t$ extra degrees of freedom to eliminate the time delay term as described by Eq. 1–3 in [2]. In an analogous way, one can easily derive φ for any non-autonomous system.

5 Conclusions

Sensitivity analysis of chaotic dynamical flows is full of mathematical and algorithmic challenges. The linear response theory, especially Ruelle's formalism, allows us to better understand how different dynamical features of a system affect its sensitivity. In particular, we can rigorously decompose the linear response formula into three separate ingredients: unstable, neutral, and stable. This concept has been utilized in recently developed algorithms such as the space-split sensitivity (S3). The unstable part represents the effect of the SRB measure gradient, which requires computing second derivatives of coordinate charts describing unstable manifolds and differentiating Lyapunov vectors in all unstable directions. The neutral and stable parts, as their names suggest, reflect the contributions of the parametric perturbation along the center (tangent to the flow) and stable manifolds, respectively. In general, any of these three terms might significantly contribute to the total linear response. The example of Lorenz 63 clearly indicates that neglecting the unstable or neutral term leads to large errors.

Despite their elegance, rigor and accuracy, direct linear response algorithms have certain flaws. First of all, they are expensive. The leading flop count may be proportional even to the cube of the number of positive Lyapunov exponents. In addition to that, the non-hyperbolic behavior of larger systems could cause numerical instabilities making the computation of measure gradients difficult. We observed that the most expansive components of the measure gradient tend to be significantly smaller in norm compared to the other ones. This critical observation led us to the conjecture that the unstable contribution could potentially be reduced if the effect of the larger components of the measure gradient is eliminated. To make the unstable part small, regardless of the choice of a parameter with respect to which linear response is computed, one could choose an *aligned* objective function J . We show

that if J is represented by the unstable divergence of a smooth vector field such that the directional derivative in the most expansive direction is dominant, the majority of the measure gradient components could be killed. Our experiment on the hyperchaotic *coupled sawtooth map* confirms that the unstable part can be significantly reduced through an appropriate selection of J .

While the idea of finding an *aligned* J may seem to be a purely theoretical concept, we argue that this result could be critical for practitioners as well. Indeed, spatially extended high-dimensional chaotic systems with statistical homogeneity in space do allow for different representations of J . In particular, the objective function, which typically equals the spatial average of system coordinates or higher-order moments, can be represented by an arbitrary linear combination of individual coordinate terms. Consequently, this gives us freedom in choosing J and increases the probability of finding an *aligned* J as the system's dimension grows. This conjecture is verified by eliminating the unstable and, consequently, the neutral part from the full S3 algorithm. Leaving the stable contribution alone, we accurately approximate sensitivities in both the Lorenz 96 and Kuramoto–Sivashinsky models.

Two primary goals were achieved in this work. First, we presented the full linear response algorithm with critical analysis of its major parts and potential applications. Second, based on our analysis, we proposed a reduced variant of S3 that has been shown to be sufficient for some higher-dimensional systems. Our results indicate that, in systems with statistical homogeneity, sensitivities could be accurately approximated by projecting out the unstable components from the tangent solution. Hence, the effect of the SRB measure change can be negligible for a wide range of parameters. We showed that when the Lyapunov spectrum collapses, which typically happens when the system moves from a non-chaotic to chaotic regime, the stable term alone is not enough. Our future work shall investigate how likely this scenario is in real-world engineering applications. If this is a rare event, further developments of well-established shadowing methods would not be necessary. Otherwise, one could consider extracting some parts of the unstable contribution to correct the reduced algorithm.

Acknowledgements The authors acknowledge the MIT SuperCloud and Lincoln Laboratory Supercomputing Center for pro-

viding HPC resources that have contributed to the research results reported within this paper

Author contributions AAS has done conceptualization, methodology, software, validation, formal analysis, investigation, data curation, writing—original draft, writing—review and editing, and visualization; QW did conceptualization, supervision, and funding acquisition.

Funding Open Access funding provided by the MIT Libraries. This work was funded by U.S. Department of Energy Grant No. DE-NA-0003993.

Data availability The datasets used for this manuscript are not publicly available, but they can be generated using the attached scripts. They are also available from the corresponding author on request.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A Full space-split algorithm—description, pseudocode and complexity analysis

The purpose of this section is to extend the discrete version of S3 [40] to continuous chaos and present the structure of the full linear response algorithm. We rely on the three-term splitting defined by Eq. 8. The major difference between the discrete and continuous variants of S3 is that, in the latter, we additionally project out the neutral component from the regularized tangent solution v . The computation of the stable part involves solving a linear system for $c^i, i = 0, 1, \dots, m$, because the vector tangent to the center subspace, f , is generally not orthogonal to the basis of the expanding subspace. That linear system is derived in Sect. 2.1. Another consequence of the three-term splitting is the emergence of the neutral contribution of the linear response. For-

tunately, as shown in Eqs. 15–19, this part of the algorithm re-uses some ingredients of the stable contribution and only requires computing $K \in \mathbb{Z}^+$ k -time correlations through ergodic-averaging. Finally, the evaluation of the unstable part also requires some adjustments. Equation 21 indicates that we need $c^i, i = 0, 1, \dots, m$, their unstable derivatives b , and derivatives of the SRB measure represented by g . We acknowledge that the computation of the SRB measure gradient is agnostic to the presence of the center manifold. Using the measure preservation property and chain rule on smooth manifolds, one can derive exponentially converging recursive formulas for g . The reader is referred to the authors’ previous work published in [41] for a detailed derivation and analysis of a trajectory-driven algorithm for g . Therefore, we only need to modify the way b is computed in the presence of the neutral subspace. Once b is found, the unstable part is computed similarly to its neutral counterpart, by summing up K k -time correlations.

Note that $b^{i,j}$ is defined as the directional derivative of c^i computed along the j -th basis vector q^j . While the regularized form of the unstable contribution (RHS of Eq. 21) involves only self-derivatives of c^i , i.e., $b^{i,i}$ with $i = j$, we show that in order to find a trajectory-following recursion, we also need all possible cross-derivatives of c^i . The main tool used in the derivation of these formulas is the measure-based parameterization of local unstable manifolds with orthonormal gradients [41]. It means that the m -dimensional unstable manifold U_k including x_k , i.e., the point of M crossed by the trajectory at the k -th time step, is parameterized as follows: $x_k(\xi) : [0, 1]^m \rightarrow U_k \subset M$ such that $x_k(\xi)$ is the multivariate inverse cumulative distribution (quantile function) and $\nabla_{\xi_k} x_k = Q_k$. In this context, the marginal SRB density ρ_k defined on U_k can be viewed as the probability density function (PDF) of the uniform measure nonlinearly re-distributed by $x_k(\xi)$. The chart coordinates ξ_k are updated step by step to ensure the orthogonality of the gradient $\nabla_{\xi_k} x_k = [\partial_{\xi_k^1} x_k, \dots, \partial_{\xi_k^m} x_k]$. A more rigorous description and analysis of this coordinate transformation can be found in [41].

To obtain $b^{i,j}, i = 0, 1, \dots, m, j = 1, \dots, m$, we simply differentiate Eq. 10, Eq. 12 and the constraint $v \cdot f = 0$ with respect to all components of ξ , apply the chain rule, and solve a linear system with $m(m + 1)$ equations and the same number of unknowns. Notice that, assuming $\nabla_{\xi_k} x_k = Q_k$, the directional derivatives

along q^i are the same as parametric derivatives with respect to ξ^i .

Differentiation of Eq. 10 with respect to ξ_{k+1}^j yields

$$\begin{aligned} \partial_{\xi_{k+1}^j} v_{k+1} &:= w_{k+1}^j = \\ \partial_{\xi_{k+1}^j} r_{k+1} &- \sum_{l=1}^m b_{k+1}^{l,j} q_{k+1}^l + c_{k+1}^l p^{l,j} \\ &- b_{k+1}^{0,j} f_{k+1} - c_{k+1}^0 Df_{k+1} q_{k+1}^j, \end{aligned} \tag{A1}$$

where $p^{i,j} := \partial_{\xi_j} q^i$. In the above equation, we used the following identity,

$$\partial_{\xi_{k+1}^j} f_{k+1} = Df_{k+1} \partial_{\xi_{k+1}^j} x_{k+1} = Df_{k+1} q_{k+1}^j.$$

Consequently, differentiating Eq. 12, i.e., constraint enforcing $v \cdot q = 0$, with respect to ξ_{k+1}^j gives

$$\begin{aligned} b_{k+1}^{i,j} &= p_{k+1}^{i,j} \cdot (r_{k+1} - c_{k+1}^0 f_{k+1}) + q_{k+1}^i \cdot \partial_{\xi_{k+1}^j} r_{k+1} \\ &- b_{k+1}^{0,j} q_{k+1}^i \cdot f_{k+1} - c_{k+1}^0 q_{k+1}^i \cdot Df_{k+1} q_{k+1}^j. \end{aligned} \tag{A2}$$

To eliminate w from the linear system, we differentiate the constraint $v \cdot f = 0$ with respect to ξ_{k+1}^j and plug Eq. A1 to obtain

$$\begin{aligned} -v_{k+1} \cdot Df_{k+1} q_{k+1}^j &= w_{k+1}^j \cdot f_{k+1} = \\ \partial_{\xi_{k+1}^j} r_{k+1} \cdot f_{k+1} & \\ - \sum_{l=1}^m b_{k+1}^{l,j} q_{k+1}^l \cdot f_{k+1} &+ c_{k+1}^l p^{l,j} \cdot f_{k+1} \\ - b_{k+1}^{0,j} f_{k+1} \cdot f_{k+1} &- c_{k+1}^0 Df_{k+1} q_{k+1}^j \cdot f_{k+1}. \end{aligned} \tag{A3}$$

Finally, by combining Eq. A2–A3, we derive the following linear system for $b^{i,j}$, $i = 0, 1, \dots, m$, $j = 1, \dots, m$,

$$\begin{aligned} (f_{k+1} \cdot f_{k+1}) b_{k+1}^{0,j} &+ \sum_{l=1}^m (q_{k+1}^l \cdot f_{k+1}) b_{k+1}^{l,j} = d_{k+1}^{0,j}, \\ (q_{k+1}^i \cdot f_{k+1}) b_{k+1}^{0,j} &+ b_{k+1}^{i,j} = d_{k+1}^{i,j}, \quad i, j = 1, \dots, m, \end{aligned} \tag{A4}$$

where

$$\begin{aligned} d_{k+1}^{0,j} &:= v_{k+1} \cdot Df_{k+1} q_{k+1}^j + \partial_{\xi_{k+1}^j} r_{k+1} \cdot f_{k+1} \\ &- \sum_{l=1}^m c_{k+1}^l p_{k+1}^{l,j} \cdot f_{k+1} - c_{k+1}^0 Df_{k+1} q_{k+1}^j \cdot f_{k+1}, \\ d_{k+1}^{i,j} &:= p_{k+1}^{i,j} \cdot (r_{k+1} - c_{k+1}^0 f_{k+1}) + q_{k+1}^i \cdot \partial_{\xi_{k+1}^j} r_{k+1} \\ &- c_{k+1}^0 q_{k+1}^i \cdot Df_{k+1} q_{k+1}^j, \quad i, j = 1, \dots, m. \end{aligned} \tag{A5}$$

The Schur complement of System A4–A5 consists of m^2 constant-diagonal blocks. Their values are exactly the same as the corresponding entries of S . Therefore, if the inverse S^{-1} is available, we can directly compute the sought-after quantities,

$$\begin{aligned} b_{k+1}^{i,j} &= (S_{k+1}^{-1})^{i:} \cdot d_{k+1}^{1:m,j} - \sum_{l=1}^m (S_{k+1}^{-1})^{il} \frac{q_{k+1}^l \cdot f_{k+1}}{f_{k+1} \cdot f_{k+1}} d_{k+1}^{0,j} \\ &= (S_{k+1}^{-1})^{i:} \cdot \left(d_{k+1}^{1:m,j} - \frac{d_{k+1}^{0,j}}{f_{k+1} \cdot f_{k+1}} Q_{k+1}^T f_{k+1} \right), \\ &i, j = 1, \dots, m, \end{aligned} \tag{A6}$$

where $(S^{-1})^{ij}$ indicates the entry of S^{-1} corresponding to its i -th row and j -th column. Analogously, $d^{1:m,j}$ denotes the m -dimensional array including all $d^{i,j}$ for all $i = 1, \dots, m$ and a fixed j . Once $b^{i,j}$ for all $i, j = 1, \dots, m$ is computed, $b^{0,j}$ and w^j , $j = 1, \dots, m$ can be evaluated directly using Eq. A1 and Eq. A4.

Based on Eq. A1–A6, we can now construct a trajectory-following iteration to compute b . These equations involve some ingredients previously derived for the stable and neutral parts. The new quantities are the parametric derivatives of the basis vectors p , i.e., derivatives of Lyapunov vectors, and $\partial_{\xi_{k+1}^j} r_{k+1}$. The former are computed using the procedure for g extended by an extra low-cost projection [40]. Using the definition of r_{k+1} and all underlying quantities, we apply the chain rule to expand $\partial_{\xi_{k+1}^j} r_{k+1}$,

$$\partial_{\xi_k^j} r_{k+1} = D^2\varphi(v_k, q_k^j) + D\varphi_k w_k^j + D\partial_s\varphi_k q_k^j, \tag{A7}$$

where $D^2\varphi(a, b)$ denotes the second-order bilinear form whose i -th component equals $(D^2\varphi(a, b))^i = \partial_{x^k} \partial_{x^l} \varphi^i a^k b^l$ (per Einstein’s summation convention), while $D\partial_s\varphi$ denotes the phase-space Jacobian of parametric derivative of φ . Note also that Eq. A7 needs to be further re-scaled by the Jacobian of the coordinate transformation from ξ_k to ξ_{k+1} . Without loss of generality, we can choose $\xi = 0$ and show that the Jacobian of coordinate transformation is a by-product of the iterative algorithm for the basis vectors q [41]. Based on the above derivations, Sects. 2.1 and [40], Algorithm 2 summarizes all the steps required to approximate the full linear response of a hyperbolic flow. While the most important aspects are covered in this work, the reader is referred to these two external references for a rigorous justification of all other parts.

Algorithm 2 Space-split sensitivity algorithm for hyperbolic flows

Input: N, K, T, n, m
Output: $d\langle J \rangle / ds \approx (s + c + u) / N$

- 1: Randomly generate: $x_0, v_0, Q_0, a_0^{i,j}, w_0^i$ for all $i, j = 1, \dots, m$
- 2: Set $s = c = u = 0$
- 3: **for** $k = 0, \dots, N - 1$ **do**
- 4: **if** $k \geq T$ **then**
- 5: $s := s + DJ_k \cdot v_k$
- 6: $u := u - J_k (u_k + u_{k-1} + \dots + u_{k-K+1})$
- 7: $c := c + DJ_k \cdot f_k (c_k^0 + c_{k-1}^0 + \dots + c_{k-K+1}^0)$
- 8: **end if**
- 9: $P_{k+1} = D\varphi_k Q_k$
- 10: QR-factorize P_{k+1} : $Q_{k+1} R_{k+1} = P_{k+1}$
- 11: Find the inverse of R_{k+1}
- 12: **for** $i = 1, \dots, m, j = 1, \dots, i$ **do** $\tilde{a}_{k+1}^{i,j} = D^2\varphi_k(q_k^i, q_k^j) + D\varphi_k a_k^{i,j}$
- 13: **end for**
- 14: **for** $i = 1, \dots, m, j = 1, \dots, i$ **do** $a_{k+1}^{i,j} = \tilde{a}_{k+1}^{p,q} (R^{-1})_{k+1}^{pi} (R^{-1})_{k+1}^{qj}$
- 15: **end for**
- 16: **for** $i = 1, \dots, m$ **do**
- 17: **for** $p, q = 1, \dots, m$ **do** $(\partial_{\xi_{k+1}^i} R_{k+1})^{pq} = \begin{cases} q_{k+1}^p \cdot a_{k+1}^{p,i}, & \text{if } p = q \\ q_{k+1}^p \cdot a_{k+1}^{q,i} + q_{k+1}^q \cdot a_{k+1}^{p,i}, & \text{if } p < q \\ 0, & \text{otherwise} \end{cases}$
- 18: **end for**
- 19: $g_{k+1}^i = -\text{tr}(\partial_{\xi_{k+1}^i} R_{k+1})$
- 20: **end for**
- 21: **for** $i, j = 1, \dots, m$ **do** $p_{k+1}^{i,j} = a_{k+1}^{i,j} - q_{k+1}^l (\partial_{\xi_{k+1}^j} R_{k+1})^{li}$
- 22: **end for**
- 23: $S_{k+1} = I - Q_{k+1}^T f_{k+1} (Q_{k+1}^T f_{k+1})^T / f_{k+1} \cdot f_{k+1}$
- 24: Find the inverse of S_{k+1}
- 25: $r_{k+1} = D\varphi_k v_k + \chi_{k+1}$
- 26: $z_{k+1} = Q_{k+1}^T (r_{k+1} - (f_{k+1} \cdot r_{k+1}) / (f_{k+1} \cdot f_{k+1}) f_{k+1})$
- 27: **for** $i = 1, \dots, m$ **do** $c_{k+1}^i = (S_{k+1}^{-1})^{ij} z_{k+1}^j$
- 28: **end for**
- 29: $c_{k+1}^0 = f_{k+1} \cdot (r_{k+1} - c_{k+1}^i q_{k+1}^i) / f_{k+1} \cdot f_{k+1}$
- 30: $v_{k+1} = r_{k+1} - c_{k+1}^i q_{k+1}^i - c_{k+1}^0 f_{k+1}$
- 31: **for** $i = 1, \dots, m$ **do** $\partial_{\xi_k^i} r_{k+1} = D^2\varphi_k(v_k, q_k^i) + D\varphi_k w_k^i + D\partial_s\varphi_k q_k^i$
- 32: **end for**
- 33: $\nabla_{\xi_{k+1}} r_{k+1} = \nabla_{\xi_k} r_{k+1} R_{k+1}^{-1}$
- 34: **for** $i = 1, \dots, m$ **do**
- 35: $d_{k+1}^{0,i} = v_{k+1} \cdot Df_{k+1} q_{k+1}^i + \partial_{\xi_{k+1}^i} r_{k+1} \cdot f_{k+1} - c_{k+1}^l p_{k+1}^{l,i} \cdot f_{k+1} - c_{k+1}^0 Df_{k+1} q_{k+1}^i \cdot f_{k+1}$
- 36: **end for**
- 37: **for** $i, j = 1, \dots, m$ **do**
- 38: $d_{k+1}^{i,j} = p_{k+1}^{i,j} \cdot (r_{k+1} - c_{k+1}^0 f_{k+1}) + q_{k+1}^i \cdot \partial_{\xi_{k+1}^j} r_{k+1} - c_{k+1}^0 q_{k+1}^i \cdot Df_{k+1} q_{k+1}^j$
- 39: **end for**
- 40: **for** $i, j = 1, \dots, m$ **do** $b_{k+1}^{i,j} = (S_{k+1}^{-1})^{i,:} \cdot (d_{k+1}^{1:m,j} - d_{k+1}^{0,j} / (f_{k+1} \cdot f_{k+1}) Q_{k+1}^T f_{k+1})$
- 41: **end for**
- 42: **for** $i = 1, \dots, m$ **do** $b_{k+1}^{0,i} = 1 / (f_{k+1} \cdot f_{k+1}) (d_{k+1}^{0,i} - (q_{k+1}^l \cdot f_{k+1}) b_{k+1}^{l,i})$
- 43: **end for**
- 44: **for** $i = 1, \dots, m$ **do** $w_{k+1}^i = \partial_{\xi_{k+1}^i} r_k - b_{k+1}^{l,i} q_{k+1}^l - c_{k+1}^l p_{k+1}^{l,i} - b_{k+1}^{0,i} f_{k+1} - c_{k+1}^0 Df_{k+1} q_{k+1}^i$
- 45: **end for**
- 46: Save the two scalars: $u_{k+1} = b_{k+1}^{i,i} + c_{k+1}^i g_{k+1}^i$ and c_{k+1}^0
- 47: Advance the iteration: $x_{k+1} = \varphi(x_k)$
- 48: **end for**

Update stable (s), neutral (c) and unstable (u) contributions [7,40].

Update basis vectors Q and transformation Jacobian R . See [13] for derivation and convergence analysis of Lyapunov basis.

Compute all m components of g . This part requires solving $m^2 2^{nd}$ -order tangent equations. Run once per time step, regardless of $\dim(J)$ and $\dim(s)$. See [41] for original derivation.

Compute derivatives of Lyapunov vectors [40].

Modification of the general discrete S3 from [40]. This part computes c, b, w , according to the derivation presented in this section and Section 2.1.

The input parameter T is to allow all the recursions to converge before the linear response contributions are collected. Note that Algorithm 2 is agnostic to the time integration method, which directly affects φ and hence the cost of computing its derivatives. In Appendix B, we derive relevant differentiation operators for the midpoint scheme.

Assuming both the objective function J and parameter s are scalars, the computational cost of Algorithm 2 depends on three parameters: the trajectory length N , dimension of both the system n and unstable subspace m . In this case, the most expensive part is the computation of the SRB density gradient (Lines 12-20). This chunk of the algorithm solves m^2 second-order tangent equations (Line 12) and performs double contraction against the transformation Jacobian (Line 13) to stabilize the iteration, which costs $\mathcal{O}(n^3 m^2 + n m^3)$ floating point operations (flops) per time step. If s is an n_s -dimensional vector, then the majority of the modified part of Algorithm 2 (Lines 23-45) will need to be repeated n_s times, which costs $\mathcal{O}(n_s (n^3 m + m^2 n))$ flops per time step. Finally, Lines 4-8 would need to be repeated n_J times if J was an n_J -dimensional vector. This would incur an extra cost proportional to $\mathcal{O}(n_J n_s n)$ flops. Therefore, assuming $\max(m, n_s, n_J) \ll n$, the leading flop count term of the total cost of Algorithm 2 is

$$\mathcal{O}\left(N n^3 (m^2 + n_s m)\right). \tag{A8}$$

Note that the most important factor in determining the total cost is the system's dimension n . This number is cubed because of the contraction of the second-order operator with two different vectors (Line 12). In practice, however, the linear differentiation operators (Jacobians, Hessians) have sparse/banded structure. This usually happens in case of PDE-related dynamical systems that have been derived using standard discretization methods such as the finite element method. The major consequence of the local structure is that the cost of evaluating first- and second-order operator-vector contractions is in fact linear to the dimension of the system. Therefore, the leading term of the flop count dramatically decreases to

$$\mathcal{O}\left(N n (m^3 + n_s m^2 + n_s n_J)\right). \tag{A9}$$

Appendix B Computing derivatives of the time integrator and implicit schemes

Both Algorithm 1 and Algorithm 2 require computing first-order derivatives in phase space as well as parametric derivatives of φ . The latter also requires second-order derivatives to compute g and w . They are products of the chain rule applied to the discrete version of the time-continuous system. The computational cost of evaluating these quantities heavily depends on the time integrator. For the Euler method, for example, differentiation of φ is equally expensive as differentiation of f . In this paper, we use second- and fourth-order fully-explicit Runge–Kutta schemes, which involve nested functions. If the system is sparse and its dimension n is large, it is efficient to compute all the tensor-vector contractions as we go rather than evaluating and storing large Jacobians and Hessians. Therefore, our aim is to use the chain rule to express all contraction types appearing in both algorithms such as $D\varphi v$ in terms of similar tensor-vector products involving derivatives of f only. In this section, we present derivations for the second-order Runge–Kutta map defined by Eq. 4. Analogous expressions for the fourth-order scheme can be found in the attached Python code.

For the midpoint method, $\varphi(x_k)$ is defined as

$$\varphi(x_k) = x_k + \Delta t f\left(x_k + \frac{\Delta t}{2} f(x_k)\right) \tag{B10}$$

$$:= x_k + \Delta t f(x_p) = x_k + \Delta t f_p,$$

where $x_p := x_k + \Delta t/2 f(x_k)$. Therefore, for any vector $v \in \mathbb{R}^n$,

$$D\varphi_k v = v + \Delta t Df_p v + \frac{\Delta t^2}{2} Df_p Df_k v, \tag{B11}$$

with $Df_k = Df(x_k)$ and $Df_p = Df(x_p)$, in compliance with our notation convention. Differentiating Eq. B11 once more and contracting it against yet another vector $a \in \mathbb{R}^n$, we obtain the following relation,

$$\begin{aligned} D^2\varphi_k(v, a) = & \Delta t D^2 f_p \left(v + \frac{\Delta t}{2} Df_k v, a \right) + \\ & \frac{\Delta t^2}{2} D^2 f_p \left(v + \frac{\Delta t}{2} Df_k v, Df_k a \right) + \\ & \frac{\Delta t^2}{2} Df_p D^2 f_k(v, a). \end{aligned} \tag{B12}$$

Recall that $D^2\varphi_k(v, a) \in \mathbb{R}^n$. Assuming f also depends on a scalar parameter s , the parametric deriva-

tive of Eq. B10 expands as follows,

$$\partial_s \varphi_k = \chi_{k+1} = \Delta t \partial_s f_p + \frac{\Delta t^2}{2} Df_p \partial_s f_k, \tag{B13}$$

where $\partial_s f_k = \partial f / \partial s(x_k)$. The final relevant contraction, $D\partial_s \varphi_k v$, involves mixed parametric and phase-space derivatives and is obtained by differentiating Eq. B13,

$$\begin{aligned} D\partial_s \varphi_k v = & \Delta t D\partial_s f_p \left(v + \frac{\Delta t}{2} Df_k v \right) + \\ & \frac{\Delta t^2}{2} D^2 f_p \left(v + \frac{\Delta t}{2} Df_k v, \partial_s f_k \right) + \\ & \frac{\Delta t^2}{2} Df_p D\partial_s f v. \end{aligned} \tag{B14}$$

We highlight the fact that, for the midpoint method, each tensor-vector product involving φ requires the evaluation of $\mathcal{O}(1)$ similar products containing f . The fourth-order Runge–Kutta scheme is in fact a four-level nested map from x_k to x_{k+1} . In this case, the Hessian-vector contraction requires about 20 such evaluations. For sparse systems, however, the cost of a single evaluation of $Df v$, $D^2 f(a, v)$, $D\partial_s f v$ is linear in n .

An implicit scheme is a common choice for stiff systems. That choice does not affect our linear response algorithms. The only part that needs to be modified is the way the products appearing in Eq. B10–B14 are computed. Let us consider a generic implicit scheme,

$$h(x_k, x_{k+1}) = 0, \tag{B15}$$

where $x_{k+1} = \varphi(x_k)$. Assuming x_k is known, the n -dimensional nonlinear system defined by Eq. B15 is typically solved for x_{k+1} using a standard solver such as the Newton–Raphson method. Differentiating Eq. B15 with respect to x_k and multiplying both sides by a vector v , we obtain the following system,

$$\frac{\partial h}{\partial x_{k+1}} D\varphi_k v = -\frac{\partial h}{\partial x_k} v, \tag{B16}$$

where $\partial h / \partial x_k$ and $\partial h / \partial x_{k+1}$ are the $n \times n$ Jacobian matrices of h with respect to x_k and x_{k+1} , respectively, both evaluated at (x_k, x_{k+1}) . If both x_k and x_{k+1} are known, the linear system defined by Eq. B16 can be solved for $D\varphi_k v$, which is a necessary ingredient of our linear response algorithms. To compute other tensor-vector products, we further differentiate Eq. B16, apply the chain rule as presented above, and formulate analogous linear systems.

References

1. Abramov, R.V., Majda, A.J.: Blended response algorithms for linear fluctuation-dissipation for complex nonlinear dynamical systems. *Nonlinearity* (2007). <https://doi.org/10.1088/0951-7715/20/12/004>
2. Arecchi, F.T., Giacomelli, G., Lapucci, A., et al.: Two-dimensional representation of a delayed dynamical system. *Phys. Rev. A* (1992). <https://doi.org/10.1103/PhysRevA.45.R4225>
3. Arnold, L.: *Random Dynamical Systems*. In: *The multiplicative ergodic theorem on bundles and manifolds*. Springer, Berlin (1998). https://doi.org/10.1007/978-3-662-12878-7_4
4. Benettin, G., Galgani, L., Giorgilli, A., et al.: Lyapunov characteristic exponents for smooth dynamical systems and for hamiltonian systems; a method for computing all of them. Part 2. *Numer. Appl. Mecc.* **15**, 21–30 (1980). <https://doi.org/10.1007/BF02128237>
5. Blonigan, P.: *Least squares shadowing for sensitivity analysis of large chaotic systems and fluid flows*. PhD thesis, Massachusetts Institute of Technology (2016)
6. Blonigan, P.J., Wang, Q.: Least squares shadowing sensitivity analysis of a modified Kuramoto-Sivashinsky equation. *Chaos, Solitons Fractals* **64**, 16–25 (2014). <https://doi.org/10.1016/j.chaos.2014.03.005>
7. Chandramoorthy, N.: *An efficient algorithm for sensitivity analysis of chaotic systems*. PhD thesis, Massachusetts Institute of Technology (2021)
8. Chandramoorthy, N., Wang, Q.: On the probability of finding a nonphysical solution through shadowing. *J. Comput. Phys.* (2021). <https://doi.org/10.1016/j.jcp.2021.110389>
9. Chandramoorthy, N., Wang, Q.: Efficient computation of linear response of chaotic attractors with one-dimensional unstable manifolds. *SIAM J. Appl. Dyn. Syst.* (2022). <https://doi.org/10.1137/21M1405599>
10. Chandramoorthy, N., Fernandez, P., Talnikar, C., et al.: Feasibility analysis of ensemble sensitivity computation in turbulent flows. *AIAA J.* **57**(10), 4514–4526 (2019). <https://doi.org/10.2514/1.J058127>
11. Chernov, N.I.: Limit theorems and Markov approximations for chaotic dynamical systems. *Probab. Theory Relat. Fields* **101**, 321–362 (1995). <https://doi.org/10.1007/BF01200500>
12. De Nittis, G., Lein, M.: *Linear response theory: an analytic-algebraic approach*. Springer, Cham (2017). <https://doi.org/10.1007/978-3-319-56732-7>
13. Ershov, V.E., Potapov, A.B.: On the concept of stationary Lyapunov basis. *Physica D* **118**, 167–198 (1998). [https://doi.org/10.1016/S0167-2789\(98\)00013-X](https://doi.org/10.1016/S0167-2789(98)00013-X)
14. Galavotti, G., Cohen, E.G.D.: Dynamical ensembles in stationary states. *J. Stat. Phys.* **80**, 931–970 (1995). <https://doi.org/10.1007/BF02179860>
15. Giacomelli, G., Meucci, R., Politi, A., et al.: Defects and spacelike properties of delayed dynamical systems. *Phys. Rev. Lett.* (1994). <https://doi.org/10.1103/PhysRevLett.73.1099>
16. Haskey, S.R., Lanctot, M.J., Liu, Y.Q., et al.: Effects of resistivity and rotation on the linear plasma response to non-axisymmetric magnetic perturbations on dIII-d. *Plasma*

- Phys. Control. Fusion (2015). <https://doi.org/10.1088/0741-3335/57/2/025015>
17. Karimi, A., Paul, M.R.: Extensive chaos in the Lorenz-96 model. *Chaos* **20**(043), 105 (2010). <https://doi.org/10.1063/1.3496397>
 18. Kontani, H., Yamakawa, Y.: Linear response theory for shear modulus C_{66} and raman quadrupole susceptibility: Evidence for nematic orbital fluctuations in fe-based superconductors. *Phys. Rev. Lett.* **113**(047), 001 (2014). <https://doi.org/10.1103/PhysRevLett.113.047001>
 19. Kubo, R.: The fluctuation-dissipation theorem. *Rep. Prog. Phys.* (1966). <https://doi.org/10.1088/0034-4885/29/1/306>
 20. Kuptsov, P.V.: Fast numerical test of hyperbolic chaos. *Phys. Rev. E* **85**(015), 203 (2012). <https://doi.org/10.1103/PhysRevE.85.015203>
 21. Kuramoto, Y., Tsuzuki, T.: Persistent propagation of concentration waves in dissipative media far from thermal equilibrium. *Prog. Theor. Phys.* **55**, 356–369 (1976). <https://doi.org/10.1143/PTP.55.356>
 22. Kuznetsov, P.K.: *Hyperbolic chaos: a physicist's view*. Springer, Berlin, Heidelberg (2012). <https://doi.org/10.1007/978-3-642-23666-2>
 23. Larsson, J., Wang, Q.: The prospect of using large eddy and detached eddy simulations in engineering design, and the research required to get there. *Philosophical Trans. R. Soc. A* **372**(20130), 329 (2014). <https://doi.org/10.1098/rsta.2013.0329>
 24. Lorenz, E.: Deterministic nonperiodic flow. *J. Atmos. Sci.* **32**(10), 2022–2026 (1963)
 25. Lorenz, E.: *Predictability - a problem partly solved*, pp. 40–58. Cambridge University Press, Cambridge (2006). <https://doi.org/10.1017/CBO9780511617652.004>
 26. Lucarini, V.: Revising and extending the linear response theory for statistical mechanical systems: evaluating observables as predictors and predictands. *J. Stat. Phys.* **173**, 1698–1721 (2018). <https://doi.org/10.1007/s10955-018-2151-5>
 27. Morales, C.A., Pacifico, M.J., Pujals, E.R.: Singular hyperbolic systems. *Proc. Am. Math. Soc.* **127**, 3393–3401 (1999)
 28. Ni, A.: Hyperbolicity, shadowing directions and sensitivity analysis of a turbulent three-dimensional flow. *J. Fluid Mech.* **863**, 644–669 (2019). <https://doi.org/10.1017/jfm.2018.986>
 29. Ni, A.: Fast linear response algorithm for differentiating stationary measures of chaos. arXiv e-prints [arXiv:2009.00595](https://arxiv.org/abs/2009.00595) (2021)
 30. Ni, A.: Approximating linear response by nonintrusive shadowing algorithms. *SIAM J. Numer. Anal.* **59**, 2843–2865 (2022). <https://doi.org/10.1137/20M1388255>
 31. Ni, A., Wang, Q.: Sensitivity analysis on chaotic dynamical systems by non-intrusive least squares shadowing (NILSS). *J. Comput. Phys.* **347**, 56–77 (2017)
 32. Pazó, D., Szendro, I.G., López, J.M., et al.: Structure of characteristic lyapunov vectors in spatiotemporal chaos. *Phys. Rev. E* **78**(016), 209 (2008). <https://doi.org/10.1103/PhysRevE.78.016209>
 33. Pilyugin, S.Y.: Shadowing in dynamical systems. In *Lecture Notes in Mathematics*, Vol. 1706, Springer-Verlag, New York (1999). <https://doi.org/10.1007/BFb0093184>
 34. Ragone, F., Lucarini, V., Lunkeit, F.: A new framework for climate sensitivity and prediction: a modelling perspective. *Clim. Dyn.* **46**, 1459–1471 (2016). <https://doi.org/10.1007/s00382-015-2657-3>
 35. Ruelle, D.: Differentiation of SRB states. *Commun. Math. Phys.* **187**, 227–241 (1997). <https://doi.org/10.1007/s002200050134>
 36. Ruelle, D.: Differentiation of SRB states: correction and complements. *Commun. Math. Phys.* **234**, 185–190 (2003). <https://doi.org/10.1007/s00220-002-0779-z>
 37. Ruelle, D.: Differentiation of SRB states for hyperbolic flows. *Ergod. Theory Dyn. Syst.* **28**, 613–631 (2008). <https://doi.org/10.1017/S0143385707000260>
 38. Sivashinsky, G.I.: Nonlinear analysis of hydrodynamic instability in laminar flames - Part I. *Deriv. Basic Equ. Acta Astronaut.* **4**, 1177–1206 (1977). [https://doi.org/10.1016/0094-5765\(77\)90096-0](https://doi.org/10.1016/0094-5765(77)90096-0)
 39. Śliwiak, A.A., Wang, Q.: Differentiating densities on smooth manifolds. *Appl. Math. Comput.* (2021). <https://doi.org/10.1016/j.amc.2021.126444>
 40. Śliwiak, A.A., Wang, Q.: Space-split algorithm for sensitivity analysis of discrete chaotic systems with unstable manifolds of arbitrary dimension. arXiv e-prints [arXiv:2109.13313](https://arxiv.org/abs/2109.13313) (2021b)
 41. Śliwiak, A.A., Wang, Q.: A trajectory-driven algorithm for differentiating SRB measures on unstable manifolds. *SIAM J. Sci. Comput.* **44**, A312–A336 (2022). <https://doi.org/10.1137/21M1431916>
 42. Śliwiak, A.A., Chandramoorthy, N., Wang, Q.: Ergodic sensitivity analysis of one-dimensional chaotic maps. *Theor. Appl. Mech. Lett.* **10**, 438–447 (2020). <https://doi.org/10.1016/j.taml.2020.01.058>
 43. Śliwiak, A.A., Chandramoorthy, N., Wang, Q.: Computational assessment of smooth and rough parameter dependence of statistics in chaotic dynamical systems. *Commun. Nonlinear Sci. Numer. Simul.* (2021). <https://doi.org/10.1016/j.cnsns.2021.105906>
 44. Takeuchi, A.T., Yang, H., Ginelli, F., et al.: Hyperbolic decoupling of tangent space and effective dimension of dissipative systems. *Phys. Rev. E* **84**(046), 214 (2011). <https://doi.org/10.1103/PhysRevE.84.046214>
 45. Van Kekem, D.L.: Dynamics of the Lorenz-96 model: bifurcations, symmetries and waves. University of Groningen Research Database https://pure.rug.nl/ws/portalfiles/portal/65106850/1_Introduction.pdf (2018)
 46. Vaupel, J.W., Yashin, A.I.: Heterogeneity's ruses: some surprising effects of selection on population dynamics. *Am. Stat.* **39**(3), 176–185 (1985). <https://doi.org/10.2307/2683925>

47. Wang, Q.: Convergence of the least squares shadowing method for computing derivative of ergodic averages. *SIAM J. Numer. Anal.* **52**, 156–170 (2014). <https://doi.org/10.1137/130917065>
48. Xu, M., Paul, M.R.: Covariant lyapunov vectors of chaotic rayleigh-bénard convection. *Phys. Rev. E* **93**(062), 208 (2016). <https://doi.org/10.1103/PhysRevE.93.062208>
49. Young, L.S.: What are srb measures, and which dynamical systems have them? *J. Stat. Phys.* **108**, 733–754 (2002). <https://doi.org/10.1023/A:1019762724717>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.