Leveraging the Linear Response Theory in Sensitivity Analysis of Chaotic Dynamical Systems and Turbulent Flows

by

Adam Andrzej Śliwiak

M.Sc., McMaster University (2017)

B.Sc., Warsaw University of Technology (2014)

Submitted to the Department of Aeronautics and Astronautics and Center for Computational Science and Engineering
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Computational Science and Engineering
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2023

Author .....
Department of Aeronautics and Astronautics
Center for Computational Science and Engineering
May 3, 2023

Certified by .....
Qiqi Wang
Associate Professor of Aeronautics and Astronautics, MIT
Thesis Supervisor and Thesis Committee Chair

Certified by .....
Semyon Dyatlov
Associate Professor of Mathematics, MIT
Thesis Committee Member

Certified by .....
Johan Larsson
Associate Professor of Mechanical Engineering, University of Maryland
Thesis Committee Member

Accepted by .....
Jonathan P. How
Richard Cockburn Maclaurin Professor of Aeronautics and Astronautics, MIT
Chair of the Department Graduate Committee

Accepted by .....
Youssef M. Marzouk
Professor of Aeronautics and Astronautics, MIT
Co-Director, Center for Computational Science and Engineering

Leveraging the Linear Response Theory in Sensitivity Analysis of Chaotic Dynamical Systems and Turbulent Flows

by

Adam Andrzej Śliwiak

Submitted to the Department of Aeronautics and Astronautics and Center for Computational Science and Engineering
on May 3, 2023, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Computational Science and Engineering

Abstract

The linear response theory (LRT) provides a set of powerful mathematical tools for the analysis of system's reactions to controllable perturbation. In applied sciences, LRT is particularly useful in approximating parametric derivatives of observables induced by a dynamical system. These derivatives, usually referred to as sensitivities, are critical components of optimization, control, numerical error estimation, risk assessment and other advanced computational methodologies. Efficient computation of sensitivities in the presence of chaos has been a major and still unresolved challenge in the field. While chaotic systems are prevalent in several fields of science and engineering, including turbulence and climate dynamics, conventional methods for sensitivity analysis are doomed to failure due to the *butterfly effect*. This inherent property of chaos means that any pair of infinitesimally close trajectories separates exponentially fast triggering serious numerical issues.

A new promising method, known as the space-split sensitivity (S3), addresses the adverse butterfly effect and has several appealing features. S3 directly stems from Ruelle's closed-form linear response formula involving Lebesgue integrals of input-output time correlations. Its linearly separable structure combined with the chain rule on smooth manifolds enables the derivation of ergodic-averaging schemes for sensitivities that rigorously converge in *uniformly hyperbolic systems*. Thus, S3 can be viewed as an LRT-based Monte Carlo method that averages data collected through regularized tangent equations along a random orbit. Despite the recent theoretical advancements, S3 in its current form is applicable to systems with one-dimensional unstable manifolds, which makes it useless for real-world models.

In this thesis, we extend the concept of space-splitting to systems of arbitrary dimension, develop generic linear response algorithms for hyperbolic dynamical systems, and demonstrate their performance using common physical models. In particular, this work offers three major contributions to the field of nonlinear dynamics. First, we propose a novel algorithm for differentiating ergodic measures induced by chaotic systems. These quantities are integral components of the S3 method and arise from

the partial integration of Ruelle's ill-conditioned expression. Our algorithm uses the concept of quantile functions to parameterize multi-dimensional unstable manifolds and computes the time evolution of measure gradients in a recursive manner. We also demonstrate that the measure gradients can be utilized as indicators of the differentiability of statistics, and might dramatically reduce the statistical-averaging error in the case of highly-oscillatory observables. Second, we blend the proposed manifold description, algorithm for measure gradients, and linear decomposition of the input perturbation, to derive a complete set of tangent equations for all by-products of the regularization process. We prove that all the recursive equations converge exponentially fast in uniformly hyperbolic systems, regardless of the choice of initial conditions. This result is used to assemble efficient one-step Monte Carlo algorithms applicable to high-dimensional discrete and continuous-time systems. Third, we argue that the effect of measure gradient could be negligible compared to the total linear response if the model is statistically homogeneous. Consequently, one could accurately approximate the sought-after sensitivity by evolving in time a single inhomogeneous tangent that is orthogonal to the unstable subspace everywhere along an orbit. This drastically reduces the computational complexity of the full algorithm.

Every major step of theoretical and algorithmic developments is corroborated by several numerical examples. They also highlight aspects of the underlying dynamical systems, e.g., ergodic measure distributions, Lyapunov spectra, spatiotemporal structures of tangent solutions, that are relevant in the context of sensitivity analysis. This thesis considers different classes of chaotic systems, including low-dimensional discrete systems (e.g., cusp map, baker's map, multi-dimensional solenoid map), ordinary differential equations (Lorenz oscillators) and partial differential equations (Kuramoto-Sivashinsky and 3D Navier-Stokes system).

Thesis Supervisor: Qiqi Wang
Title: Associate Professor of Aeronautics and Astronautics, MIT

Acknowledgments

For of Him,
and through Him,
and to Him, are all things:
to whom be glory for ever.
Amen.
– Romans 11:36

Being a PhD student at MIT was an extremely rewarding and enjoyable part of my life. Due to the outstanding rigor and large expectations, the path of every MIT student is usually full of painful challenges and struggles. Therefore, the successful conclusion of that long journey would not be possible without a number of people that supported me academically, mentally and spiritually over the last few years.

First of all, I am grateful to my advisor, Prof. Qiqi Wang, who exposed me to the staggering field of nonlinear dynamical systems. Qiqi's invaluable support, patience and belief in success of my research project led to this 300-page thesis. All those hundreds of discussions I had with Qiqi did not only improve the quality of my thesis, but also contributed to my personal growth as a research scholar.

I would like to thank the two other members of my Thesis Committee, Prof. Johan Larsson and Prof. Semyon Dyatlov, for sharing their expertise in the field of turbulence modeling and dynamical systems theory, and for mentoring me over the course of my thesis research. By the same token, I would like to thank Prof. Jaime Peraire and Prof. Adrián Lozano-Durán for reviewing my thesis and offering helpful suggestions that led to a number of improvements of this long manuscript. I am also grateful to Prof. Raúl Radovitzky, who supervised me at the early stage of my doctoral program.

Some results presented in this thesis were obtained with the help of fellow students and several other researchers. First, I would like to acknowledge the contributions of Nisha Chandramoorthy, who co-authored manuscripts included in Chapter

2 and Chapter 4. Nisha's excellent thesis work provided theoretical background and a number of starting points and research directions for my studies described in this document. Second, I would like to thank Nikhil Oberoi, who provided a python code for the turbulent channel flow. This enabled the completion of Chapter 7. Third, I would like to thank Dr. Patrick Blonigan and Dr. Francesco Rizzi, who supervised me during my Summer 2022 internship at Sandia National Laboratories. I am also grateful to several anonymous reviewers volunteering in the peer-review process. Their comments significantly improved the quality of the submitted manuscripts as well as my general understanding of the field.

Finishing this thesis would be infeasible without my family and best friends. I am particularly grateful to my dear wife, Monika Śliwiak, for her unconditional love and always being there for me. Unlike me, Monika always believed in a happy ending of this journey. She was the one who perpetually supported me at every single stage of my MIT life, in joyful and tough moments. We celebrated all success together, and also we cried together when a failure came. I am thankful to my parents (Barbara and Andrzej Śliwiak), parents-in-law (Agnieszka and Andrzej Wójcik), sisters (Agnieszka Carewicz and Aleksandra Paszkowska) and their families, sister-in-law (Elżbieta Nieciecka) and her family, for their continuous support. Finally, I would like to thank my extended family and friends from MIT, Boston area and beyond: the Bramowski family, the Kasperkiewicz family, Fr. Norbert M. Siwiński OFM Conv., Fr. Jerzy Żebrowski OFM Conv., Fr. Jarosław Szumański, Basia Gałaj, Ania Maziarz, Michał Papaj, Grzegorz Jarczok, members of ACSEL and ISN, members of Polish Club at MIT.

This research was sponsored by the following programs and institutions: MIT Department of Aeronautics and Astronautics, MIT Center for Computational Science and Engineering, Predictive Science Academic Alliance Program 3 (PSAAP-3), National Nuclear Security Administration, Air Force Office of Scientific Research, Sandia National Laboratories, MIT Lincoln Laboratory (provided HPC resources through MIT Supercloud), Lawrence Livermore National Laboratories (provided HPC resources).

Contents

List of Figures

List of Tables

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Introduction

Chaos is the score upon which the reality is written. – Henry Miller

1.1 Linear response of chaos and its applications

Linear response theory (LRT) [44] is a branch of computational science and engineering that 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 differentiably with respect to its input. Throughout this thesis, we use the phrase "linear response" interchangeably with "sensitivity", as the main purpose of LRT is in fact the sensitivity analysis of dynamical systems.

In the presence of chaos, the classical formulation of LRT is modified. The quantity of interest is usually expressed in terms a long-time average or higher-order dis-

tribution moments, rather than instantaneous values, of a certain observable J. In particular, the sought-after system's sensitivity with respect to a scalar parameter s equals

$$\frac{d\langle J\rangle}{ds} =: \frac{d}{ds}\left(\lim_{N\to\infty}\frac{1}{N}\sum_{k=0}^{N-1}J(x_k)\right), \tag{1.1}$$

where x_k denotes the system's state (snapshot) at time k. 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 [148, 150], who rigorously derived a closed-form expression for the linear response of chaos. That formula is represented in terms of a convolution of the linear response operator, which uses the information of the unperturbed system, and imposed forcing or parameter perturbation (input). The major assumption of Ruelle's derivation is *uniform hyperbolicity*, which is a mathematical idealization of chaotic behavior. In short, a chaotic system defined on n-dimensional manifold M is hyperbolic if its tangent space TM can be split into two *covariant* subspaces: unstable and stable. They are spanned by expanding and contracting directions of the tangent space and correspond to positive and negative Lyapunov exponents (LEs), respectively. Whenever asymptotic expansion and contraction rates are uniform, we speak of the strongest version of hyperbolicity, i.e., uniform hyperbolicity. The key aspect of hyperbolicity *per se* 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 [195], which contains the statistical description of the dynamics. In the case of flows (continuous-time systems), the definition of hyperbolicity also considers the existence of the center (neutral) manifold that is parallel to the evolution vector [151]. More rigorous descriptions of this fundamental assumption is provided in the following chapters.

Solid numerical evidence found in the literature clearly indicates that uniform hyperbolicity is a sufficient, but not necessary, condition for the differentiability of statistics [18, 28]. Indeed, these empirical results are consistent with *hyperbolic hy-*

pothesis of Galavotti and Cohen [59]. 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 [126, 163], where the authors argued that the long-time averages computed for a 3D turbulence model are smooth despite local non-hyperbolic behavior.

In this thesis, we focus on autonomous discrete systems,

$$x_{k+1} = \varphi(x_k; s) \tag{1.2}$$

and flows

$$\frac{dx}{dt} = f(x; s) \tag{1.3}$$

governed by the diffeomorphic map $\varphi : M \to M$ and evolution vector f, respectively. Indeed, φ can be interpreted as a time integrator of Eq. 1.3. 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,

$$x_{k+1} = \varphi(x_k) = x_k + \Delta t\, f(x_k + \frac{\Delta t}{2} f(x_k)). \tag{1.4}$$

Assuming the systems defined by Eq. 1.2–1.2 are uniformly hyperbolic, Ruelle's linear response formula is guaranteed to hold and can be expressed as follows [148, 150],

$$\frac{d\langle J\rangle}{ds} = \sum_{t=0}^{\infty} \int_M D(J \circ \varphi^t) \cdot \chi \, d\mu, \tag{1.5}$$

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

$$\int_M h(x)\, d\mu = \lim_{T\to\infty} \frac{1}{T} \int_0^T h(x(t))\, dt \approx \frac{1}{N} \sum_{k=0}^{N-1} h(x_k) \tag{1.6}$$

for any observable $h \in L^1(\mu)$ and a sufficiently large sample size N. Thus, the right-hand side (RHS) of Eq. 1.5 could potentially be approximated by computing a sufficiently long trajectory, ergodic-averaging the integrand per Eq. 1.6, and truncating the infinite series. However, note that

$$D(J \circ \varphi^t) \cdot \chi = (DJ)_t \cdot (D\varphi)_{t-1} ... D\varphi \, \chi. \tag{1.7}$$

$(DJ)_t$ 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 subspace, the norm of that product grows exponentially fast with t,

$$\| D\varphi_{t-1} \, D\varphi_{t-2} ... D\varphi \, \chi \| \sim \mathcal{O}(\exp(\lambda_1 t)) \tag{1.8}$$

with $\lambda_1 > 0$, which means the direct evaluation of the RHS of Eq. 1.5 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 [195]. Therefore, integration by parts would be possible only if χ belongs to unstable manifolds everywhere in M, which is generally not the case.

An accurate approximation of the sensitivity or linear response defined by Eq. 1.5 could further our understanding of several fields of science and engineering dealing with chaotic behavior. They include:

- Fluid sciences: turbulence theory, aerodynamics [61, 48, 79, 126, 97, 18, 13, 155],
- Molecular dynamics [80, 157, 179, 77],
- Climate dynamics [110, 102, 143, 22],

- Statistical physics [113, 49, 101],
- Plasma physics and fusion [75, 175, 54],
- Circuit theory and electromagnetism [87, 184, 108],
- Operational research and management science [183, 171, 125],
- Social sciences [67, 53].

In both non-chaotic and chaotic realms, parametric derivatives of system outputs are often used in the development of advanced algorithms for applied analysis of complex systems. In particular, sensitivities are highly-desired in gradient-based approaches for optimal design and control [86, 61, 38, 185, 81, 196, 130, 63, 55], data assimilation [145, 119, 178, 71, 30], uncertainty quantification [9, 26, 16, 180, 116], inverse problems [168, 74, 167], grid adaptation [96, 56], data-driven methods [76, 141, 58].

Computing sensitivities in non-chaotic ergodic systems is rather simple and does not require further algorithmic developments. Indeed, the RHS of Eq. 1.5 can be rigorously re-expressed in terms of a single ergodic average involving an inhomogeneous tangent solution that is bounded at all times. This is no longer the case in the presence of positive Lyapunov exponents, which requires special treatment alleviating the butterfly effect.

1.2 State of the art in the sensitivity analysis of chaotic systems

Local sensitivity analysis could simply be performed using finite differences (FD) by running two simulations with the input parameter being equal to s and $s + \delta s$ [158]. There are two major reasons why we search for a technology that is better than FD. First, notice that we need to run independent simulations at different parameter values, compute corresponding ergodic-averages, and apply a central difference scheme

with an inherent error proportional to $\mathcal{O}(\delta s^2)$. Therefore, the FD approximation of sensitivities is affected by a combination of the statistical (Monte Carlo) error and Taylor series truncation error. Intuitively, the smaller δs, the longer trajectories we need to better approximate the slope in the parametric space. The second reason is that FD is not the best option for multi-query problems common in design optimization and uncertainty quantification. Problems of this type usually require estimating thousands of sensitivities of different observables with respect to large arrays of parameters [146, 78]. In the case of FD, the total cost is proportional to the number of parameters.

Within the family of "continuous" approaches (as opposed to "discrete" FD), the most straightforward group uses basic linear perturbation techniques. In particular, they usually solve a recursive tangent equation, i.e., a linearized version of the original nonlinear system. These approaches are also referred to as *conventional methods*, which include adjoint methods [83, 27, 197]. Motivated by the optimal control theory, the adjoint methods integrate classical tangents by parts obtaining a backwards system that does not depend on parametric perturbations. Effectively, the cost of sensitivity analysis is independent from the number of parameters. They are sometimes combined with automatic differentiation [142] or "discrete" methods [137, 197]. However, as described in Section 1.3, tangent solutions of any form represent the separation of two trajectories initiated at two different, but very close to each other, initial conditions. Thus, in the presence of chaos, they grow exponentially fast in norm, which is indeed a direct consequence of the butterfly effect. This computational inconvenience was circumvented in ensemble methods relying on the assumption of ergodicity. Instead of generating a long trajectory, Eyink et al. [52] proposed computing sensitivities of several truncated-in-time trajectories and taking the average of the partial results. While this approach does not suffer from the butterfly effect and has been proven successful for some real-world chaotic systems [29, 102], large variances of the partial estimates make the ensemble methods prohibitively expensive even for medium-sized models.

Yet another popular family of methods derives from the shadowing lemma [139]

which, under the assumption of uniform hyperbolicity, guarantees the existence of a shadowing trajectory that lies within a small distance to the reference orbit for a long time, yet finite. The concept of shadowing was used to formulate an optimization problem to find the shadowing direction, which can be directly used to estimate the sensitivity of discrete systems [181, 182]. Least-Squares Shadowing (LSS) and its derivatives were successfully applied to various low- and large-dimensional ODE and PDE systems [129, 20, 84, 154], including 3D turbulence models [126, 18]. An alternative approach utilizing the shadowing lemma was proposed in [99]. This work replaces the optimization problem with a periodic tangent/adjoint equation. In certain cases, however, shadowing solutions might be unphysical and feature a dramatically different statistical behavior compared to the unperturbed orbit [34]. Empirical evidence of the unphysicality of shadowing solutions was demonstrated in the advection-dominated regime of the Kuramoto-Sivashinsky model [21]. To the best of our knowledge, no rigorous studies that quantify or bound shadowing errors due to the problem of *non-physicality* are available.

Some methods directly stem from the Fluctuation-Dissipation Theorem (FDT) [88, 2, 1, 68]. FDT is widely used in the statistical equilibrium analysis of chaotic phenomena such as turbulent flows [115] and quantum chaos [124]. In the context of sensitivity analysis, the FDT theorem asserts that one can approximate the infinite-time average of an observable as a time convolution of the external perturbation and the linear response operator that contains a statistical description of the unperturbed system. Despite their computational efficiency, FDT-based methods often require specific assumptions for the statistical behavior (e.g., Gaussian ergodic measure) to reconstruct the linear response operator, which makes them hardly generalizable.

Recent algorithmic developments rely on the regularized variant of Ruelle's expression. Indeed, as originally proposed by Ruelle in [148], 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 [195] that is absolutely continuous only on unstable manifolds, an extra step is required before partial inte-

gration 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 [148]. These terms will be respectively referred to as the unstable and stable contribution of the sensitivity. In the case of flows (continuous-time systems), the center manifold should also be taken into account in the perturbation splitting [151]. Based on this idea of regularization of Ruelle's closed-form expression, two conceptually similar methods for the linear response emerged in the past three years. Those are the fast linear response algorithm [127] and space-split sensitivity (S3) algorithm [31, 161]. 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 [160, 162, 127]. Differentiation of SRB measures, either explicit or implicit, is by far the most complicated and expensive part of both algorithms. Relying on the assumption of uniform hyperbolicity, the authors of [31] showed that space-splitting exists and is differentiable, rigorously proved the convergence of all the components of S3 for systems with one positive Lyapunov exponent, and demonstrated the convergence of ergodic averages stemming from Ruelle's expression.

All the aforementioned approaches can generically be labelled as "math-based" methods as they are constructed based solely on mathematical theories. This broad family also includes methods that directly utilize transfer operators (e.g., Frobenius-Perron or Koopman) governing the evolution of observables and/or ergodic averages of the system [19, 98, 112, 173]. However, operator-theoretic approaches scale poorly

with the system's dimension. Most of them have been applied only to low-dimensional problems so far and, to the best of our knowledge, no examples of their applications to large physical systems are available. "Physics-based" approaches, on the other hand, rely on physical knowledge and/or empirical data concerning a specific system [132, 6, 96]. These methods are typically cheaper compared to most of the "math-based" approaches; however, they are hardly generalizable and lack rigorous error estimates.

1.3 Contributions of this work

This thesis is a collection of several contributions to the field of sensitivity analysis of chaotic dynamical systems. Its core is the generalized linear response algorithm, or the space-splitting sensitivity (S3) method, which regularizes the closed-form expression for the sensitivity derived by Ruelle (Eq. 1.5). The first version of S3 was proposed in the PhD thesis of Chandramoorthy in [28]. That preliminary work included rigorous derivation and comprehensive convergence analysis of the underlying tangent equations and ergodic-averaging schemes; however, all the results so far were limited to systems with one-dimensional unstable subspaces, i.e., models with one positive Lyapunov exponent. Systems encountered in science and engineering are typically hyperchaotic, i.e., they have more than one LE, which means that the original derivation of S3 is not applicable to real-world cases. From the preliminary study of S3 and this thesis, we learned that the space-splitting algorithm deserves further developments and investigation for the following reasons:

- S3 effectively circumvents the problem of exponentially growing perturbations, i.e, the *butterfly effect*. This is possible thanks to the engineered splitting of the input perturbation and construction of numerically stable tangent recursions for the resulting quantities.

- S3 can be formulated as a one-step local-in-time Monte Carlo scheme. It means that the underlying numerical procedure involves collecting and averaging nec-

essary quantities that are generated recursively in time. These averages can be updated every time step and there is no need to save the entire history.

- Under the assumption of *uniform hyperbolicity*, all the recursions of S3 are guaranteed to convergence exponentially fast, regardless of the number of positive LEs.

- *Uniform hyperbolicity* is the only assumption used in the development. It does not rely on any other simplifications, such as specific statistical behavior or established physical constraints.

- Thanks to the separable input-output feature of Ruelle's expression, it is possible to isolate the most expensive part of the algorithm. One can compute it once and apply for any input perturbation.

The content of Chapters 2-7, containing all the main contributions of this thesis, can be classified into three major categories:

1. **Algorithm for differentiating densities implied by arbitrary diffeomorphisims and computing ergodic measure gradients.** A necessary step in regularizing Ruelle's original expression is integration by parts. Due to the presence of nonuniform distributions, the regularized formula for the sensitivity involves an extra term representing derivatives of the ergodic measure. In the case of geometrically linear unstable manifolds, computing these new terms is rather simple and can be done by differentiating the Frobenius-Perron operator. More complicated manifolds require the development of systematic and generalizable numerical tools that could be formulated as a trajectory-driven Monte Carlo algorithm. We propose the *measure-based* parameterization of smooth manifolds that uses quantile functions. This framework enables expressing any directional derivative of the density implied by smooth nonlinear transformations in terms of first- and second-order derivatives of the coordinate chart. We show that the proposed manifold description is naturally extendable to higher dimensions, and that the density gradient formula is always invariant to any

linear change of variables. Subsequently, we utilize this machinery to describe the evolution of the SRB measure gradient in chaotic systems. By applying the chain rule on smooth manifolds, we rigorously derive a set of first- and second-order tangent equations for the Lyapunov vectors and acceleration (curvature) vectors, both of which directly contribute to the sought-after SRB measure gradient. We show that all of these tangents are immune to the butterfly effect due to the step-by-step coordinate orthonormalization. Using the aforementioned concept of linear invariance, we argue that the recursive othonormalization of the chart gradients does not violate the consistency of the proposed scheme.

2. **Generalization of the space-splitting sensitivity (S3) algorithm to hyperbolic systems of arbitrary dimension.**

   We combine the idea of space-splitting with the *measure-based* parameterization of unstable tangent subspaces. First, we propose to decompose the input perturbation using an orthonormal basis of the unstable manifold. Orthonormal bases (a.k.a. Gram-Schmidt bases) are significantly cheaper compared to their covariant counterparts and they dramatically simplify the entire algorithm. Second, we rigorously derive tangent equations for all quantities arising in the regularization of Ruelle's expression. The key aspect of this derivation is the relation between the derivatives of Lyapunov vectors and acceleration vectors, i.e., second-order derivatives of the coordinate chart. Third, we rigorously prove that all the generalized tangent equations constituting the S3 algorithm converge exponentially fast in time. Finally, these recursions are assembled into a concise ready-to-use pseudocode and its complexity is carefully studied. The generalized space-splitting algorithm is applied to physically inspired hyperchaotic systems. We also explain the interplay between the truncation parameters and their impact on the convergence of the Monte Carlo averaging scheme.

3. **Simplified S3 method and its applications.**

   The unstable contribution of the space-split algorithm requires solving costly second-order tangent equations. The number of these equations scales quadrat-

ically with the dimension of the unstable subspace. Moreover, the SRB measure that we need to differentiate is not guaranteed to be smooth in large physical systems because of their non-hyperbolic behavior. Due to the high cost and potential stability issues in the computation of the unstable contribution, we try to answer a fundamental question about its relevance in the context of practicable systems. In particular, we investigate if there exist chaotic systems, where the unstable contribution is small compared to the remaining parts of the total sensitivity. We show that the effect of the SRB measure gradient could be minimized if the objective function is aligned with the most expansive direction of the tangent space. This result could be used as a powerful handle in reducing the S3 method. We demonstrate that in large systems featuring statistically homogeneous behavior in the physical space, we have freedom in representing the objective function in infinitely many ways. Indeed, if the system's dimension grows, the likelihood of the desired alignment increases as well. Since several physical systems are statistically homogeneous, we conclude that the sensitivity could be accurately approximated through the stable contribution alone. Thus, it might be sufficient to compute a single inhomogeneous tangent solution that is orthogonal to the unstable subspace everywhere along a random trajectory. We verify this remarkable conjecture by approximating sensitivities in some physical systems, including Lorenz 96, Kuramoto-Sivashinsky and compressible Navier-Stokes.

In addition to the three major contributions listed above, this work also offers some minor advancements relevant to the broad field of computational science and engineering. These contributions were used as auxiliary tools to support our major arguments. They include:

- Numerical assessment of Lebesgue-integrability of observables induced by hyperbolic systems,
- Numerical method for approximating the Hölder exponent of real-valued functions,

- Efficient Monte Carlo integration of highly-oscillatory functions,
- A computable criterion for differentiability of statistics in low-dimensional chaotic systems.

1.4 Thesis outline and corresponding journal publications

This work is structured as follows. The introduction is followed by seven self-contained chapters that have been ordered consistently with the development process of the linear response algorithm. To a large extent, the notation is consistent throughout this thesis and every exception is clearly highlighted. To facilitate the navigation, appendices are included at the end of their respective chapters, while the cumulative bibliography is attached at the very end of this document. Subsequent chapters are briefly summarized below.

Chapter 2: Offers a tutorial explaining basic concepts related to the linear response algorithm. It explains the integration by parts on unstable manifolds, concept of SRB measure gradients and ergodic-averaging process. These ideas are illustrated using simple one-dimensional chaotic maps. The content of this chapter was published in [158].

Chapter 3: Consists of two main parts. In the first one (Sections 3.1 – 3.5), the concept of manifold description through quantile functions is introduced and a method for differentiating densities implied by generic diffeomorphisms is proposed. The second part (Sections 3.6 – 3.10) extends these concepts to limiting stationary distributions and proposes a novel ergodic-averaging algorithm to differentiate SRB densities. These two parts were published in [160] and [162], respectively.

Chapter 4: Argues that the SRB measure gradient could be used to verify the

existence of the linear response in chaotic dynamical systems. It develops a novel numerical method to assess the differentiability of statics and validates the results by approximating Hölder exponents of the statistics-vs-parameter curves. The content of this chapter was published in [159].

Chapter 5: Generalizes the S3 method to multi-dimensional hyperbolic systems with an arbitrary number of positive Lyapunov exponents. It rigorously derives and analyzes the convergence of tangent equations for all the quantities arising from the perturbation splitting and partial integration of the unstable contribution. This chapter provides a pseudocode for the full space-splitting method, analyzes its complexity and presents several numerical examples. The content of this chapter was published in [161].

Chapter 6: Extends the discrete version of the S3 method to hyperbolic flows. Its major contribution is the argument of small unstable contributions in statistically homogeneous systems. Based on the concept of alignment of the objective function, it proposes a reduced version of S3. Numerical examples involving popular physical systems, such as Lorenz 96 and Kuramoto-Sivashinsky, are included. The content of this chapter was published in [163].

Chapter 7: Focuses solely on turbulent channel flows. This chapter investigates aspects of the tangent space, e.g., Lyapunov spectra, and applies the reduced S3 method to compute sensitivities of isothermal compressible channel flows.

Chapter 8: Summarizes the main conclusions of this thesis and suggests directions for future work.

We also acknowledge that the thesis' introduction (Chapter 1) includes parts of the author's relevant journal publications [158, 159, 160, 162, 161, 163].

Chapter 2

Ergodic-averaging sensitivity analysis of one-dimensional chaos

This chapter serves as a tutorial introducing and explaining fundamental concepts used throughout the thesis. Here, we focus on ergodic-averaging sensitivity analysis of simple one-dimensional expansive maps. We demonstrate how to leverage Ruelle's formula, which is based on the theory of hyperbolic dynamics, to derive a rigorously converging and immune to the *butterfly effect* recursive scheme for the linear response. The presented numerical examples indicate computational advantage of the proposed method over naïve finite difference computations. Moreover, the new method is generalizable, i.e., it does not assume that the probability distribution in phase space is of a particular type. We thoroughly explain basic mathematical tools and ingredients of the derived method, e.g., partial integration on unstable manifolds, the idea of SRB measure gradients, decay of time correlations. These concepts are fundamental in understanding the *unstable contribution* of the space-split sensitivity (S3) method, which we develop in the following chapters.

The content of this chapter was published in *Theoretical and Applied Mechanics Letters* by *Elsevier* [158].

2.1 Parameterized one-dimensional chaotic maps and their statistical dependence on parameter

In this section, we introduce two families of perturbed one-dimensional chaotic systems that can generally be expressed as

$$x_{k+1} = \varphi(x_k; s), \qquad x_0 = x_{\text{init}}, \tag{2.1}$$

where x_{init} is a given initial condition, while s denotes a scalar parameter. Let J be a scalar observable. Our quantity of interest is the infinite-time average or *ergodic* average of J,

$$\langle J \rangle := \lim_{N\to\infty} \frac{1}{N} \sum_{i=0}^{N-1} J(x_i; s). \tag{2.2}$$

In particular, we focus on the relationship between $\langle J \rangle$ and the parameter s for φ. In addition, we review the concepts of Lyapunov exponents and ergodicity through numerical illustrations on the two maps.

2.1.1 Perturbations of the sawtooth map and their Lyapunov exponents

We consider as our first example, perturbations of the sawtooth map, also known as the dyadic transformation or Bernoull shift [47], defined in the following way:

$$x_{k+1} = \varphi(x_k; s) = 2x_k + s\sin(2\pi x_k) \bmod 1, \quad x_k \in [0, 1). \tag{2.3}$$

It is a periodic map that maps $[0, 1)$ to itself. Figure 2-1 illustrates the sawtooth map for different values of the parameter s.

A natural question that arises is whether the chosen map is actually chaotic. A visible symptom of a chaotic map is its sensitiveness to initial conditions. What does it mean in practice? Consider the sawtooth map with $s = 0$ and two different but infinitesimally close initial conditions, x and $x + \delta x$. After one iteration, the two

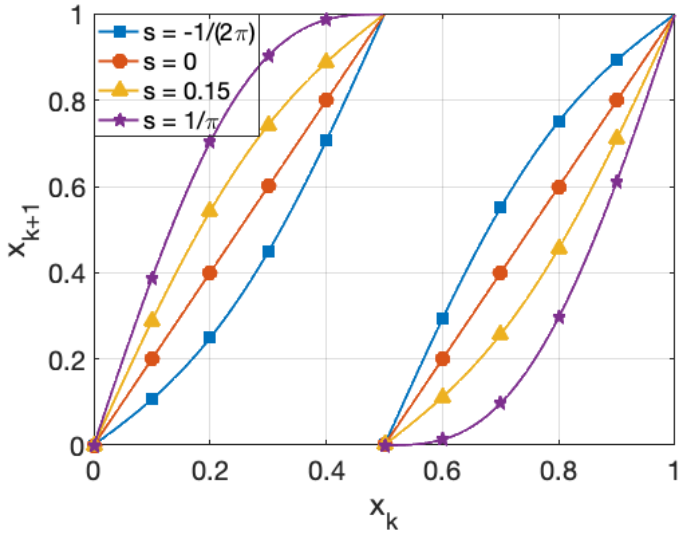

Figure 2-1: The sawtooth map at different values of parameter s. Note if $s = 0$, we obtain the classical Bernoulli shift.

trajectories will be separated by a distance of $2\,\delta x$. As we continue this recursion to the N-th iteration, that distance will rapidly grow to

$$2^N\,\delta x = \exp(\log(2)\,N)\,\delta x.$$

Therefore, we see that the two trajectories that are initially δx apart separate exponentially fast. At $s = 0$, the rate of exponential expansion is $\lambda = \log 2 \approx 0.693$. The parameter λ is known as the Lyapunov exponent (LE) and is rigorously defined as

$$\lambda(s) = \lim_{N\to\infty} \frac{1}{n} \sum_{k=0}^{N-1} \log \left| \frac{\partial \varphi}{\partial x}(x_k; s) \right|. \tag{2.4}$$

We say that a map is chaotic when its LE is positive.

Formula 2.4 requires computing the derivative of the map at points along a trajectory. Note that the value of the Lyapunov exponent does not depend on initial condition x_{init}, nor on the step k. Figure 2-2 shows that $\lambda > 0$ for all $s \in (-\frac{1}{2\pi}, \frac{1}{\pi}]$ meaning that the sawtooth map is chaotic in this regime. This can be easily justified by the observation that $\frac{\partial \varphi}{\partial x} \geq 1$ for all $x \in [0, 1]$, when s is in this regime.

We alert that simulating the sawtooth map with $s = 0$, i.e., the classical Bernoulli

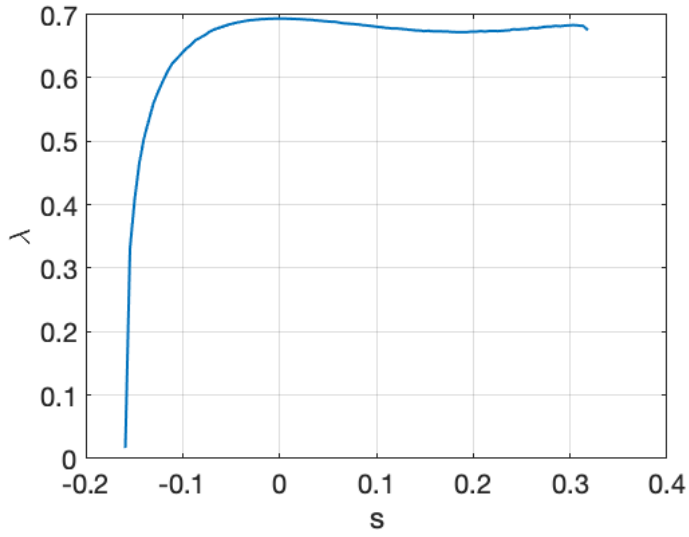

Figure 2-2: Relation between the Lyapunov exponent λ and parameter $s \in \left[-\frac{1}{2\pi}, \frac{1}{\pi}\right]$ for the sawtooth map.

shift, may require special care. Notice that all machine-representable numbers are dyadic rational, which means that they can be represented as a fraction with the denominator being a power of two. Consequently, since the Bernoulli map is simply a leftshift operation on binary digits, its numerical simulation will converge to a fixed number in most cases. More details about this problem and possible remedies can be found in Appendix 2.6.1. Note also that if $s = 0$ and x_{init} is rational, the forward orbit of x_{init} would either converge to a fixed point or be periodic, containing a finite number of distinct values within the interval $[0, 1)$. If $x_{\text{init}} = 0.1$, then all future states belong to a four-element set, $\{0.2, 0.4, 0.6, 0.8\}$, and $x_k = x_{k+4}$ for all $k > 0$. This is an example of an unstable periodic orbit; in this chapter, we are interested in chaotic orbits, which are aperiodic and unstable to perturbations.

2.1.2 A family of cusp maps and their Lyapunov exponents

Another example of a chaotic map is the cusp map $\varphi : [0, 1] \to [0, h]$ defined as follows,

$$x_{k+1} = \varphi(x_k; h, \gamma) = h - \left|\frac{1}{2} - x_k\right| - \left(h - \frac{1}{2}\right) |1 - 2x_k|^{\gamma} . \tag{2.5}$$

The above function produces a spade-shaped graph, as shown in Figure 2-3. The

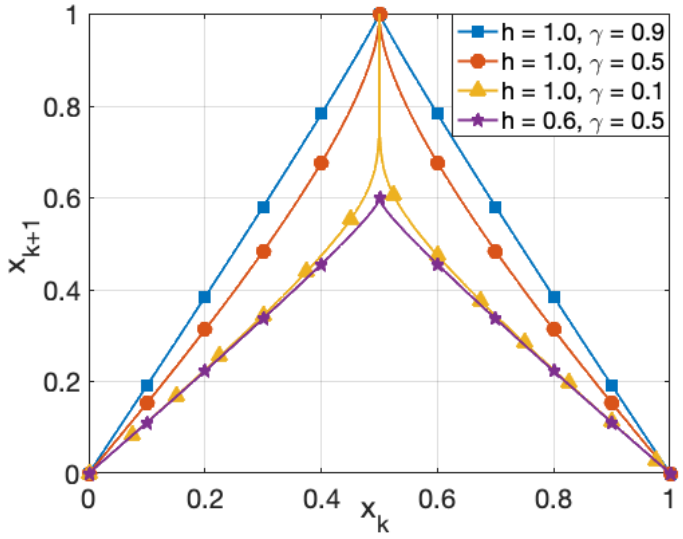

Figure 2-3: The cusp map at different values of parameters h and γ. Note all the curves include points $(0, 0)$ and $(1, 0)$, while the tip is located at $(0.5, h)$. If $h = \gamma = 1$, the map is piecewise linear, and this particular case is usually referred to as the tent map.

cusp map is a two-parameter map with $s = \{h, \gamma\}$, where h is the height, while γ is a parameter that determines the sharpness of the tip. We use the definition Eq. 4 to compute the LE of the cusp map at different values of h and γ. From the positivity of the Lyapunov exponent shown in Figure 2-4, we see that the cusp map is always chaotic if $\gamma \in [0, 1]$ and $h \geq 0.6$.

Historically, the cusp map has been used as a one-dimensional representation of the three-dimensional Lorenz'63 system [110], a set of ordinary differential equations used as a model for atmospheric convection. Specifically, the iterates of the cusp map are local maxima of the third coordinate of the Lorenz'63 system [121].

2.1.3 Statistical QoI: long-time averages

The fundamental quantity of interest (QoI) of chaotic systems that we study in this thesis is the long-time average of a certain observable. QoI is calculated in the following way:

$$\langle J \rangle \approx \frac{1}{N} \sum_{i=0}^{N-1} J(x_i), \tag{2.6}$$

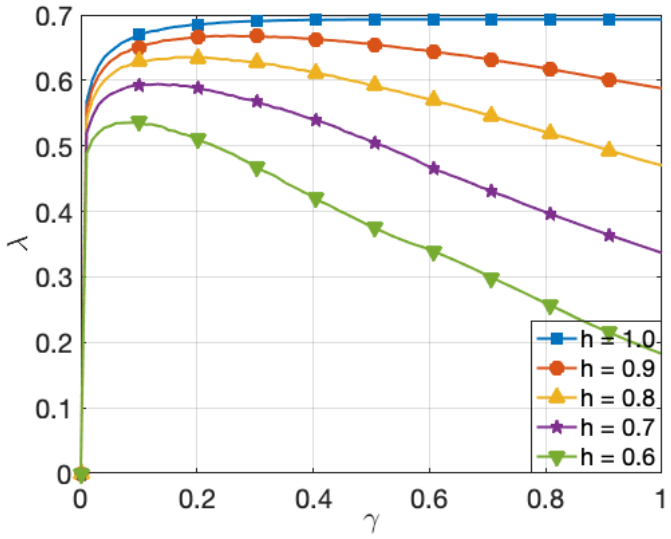

Figure 2-4: Relation between the Lyapunov exponent λ and parameter $\gamma \in [0, 1]$ for the cusp map.

where $x_{i+1} = \varphi(x_i)$. We set N to a sufficiently large number to minimize the effect of variance. Figure 2-5 illustrates examples of the mean statistics (i.e. long-time averages) and their dependence on the map parameters for both the sawtooth and cusp map. In this particular simulation, we generate 100 million samples with the initial condition chosen uniformly, at random between 0 and 1.

In the computation of long-time averages of the objective function, we used the concept of ergodicity. This property guarantees that long-time averages do not depend on the initial condition. That is, the time average of the objective function (right hand side of Eq. 2.6) converges, as $N \to \infty$, to a value independent of the initial condition x_0, for almost every x_0 chosen uniformly between 0 and 1. This limit equals the expected value of the same objective function with respect to the underlying *ergodic*, *invariant* probability distribution usually denoted by ρ. The property of invariance means that for any open interval $A \subset (0, 1)$, $\rho(A) = \rho(\varphi^{-1}(A))$. Such a probability distribution ρ is known as the Sinai-Ruelle-Bowen (SRB) density or measure [194], and it generally represents a non-uniform distribution. One of the exceptions is the sawtooth map at $s = 0$, in which case $\rho'(x)$ is zero everywhere on $(0, 1)$. A rigorous definition SRB measures and further explanations in the context of multi-dimensional

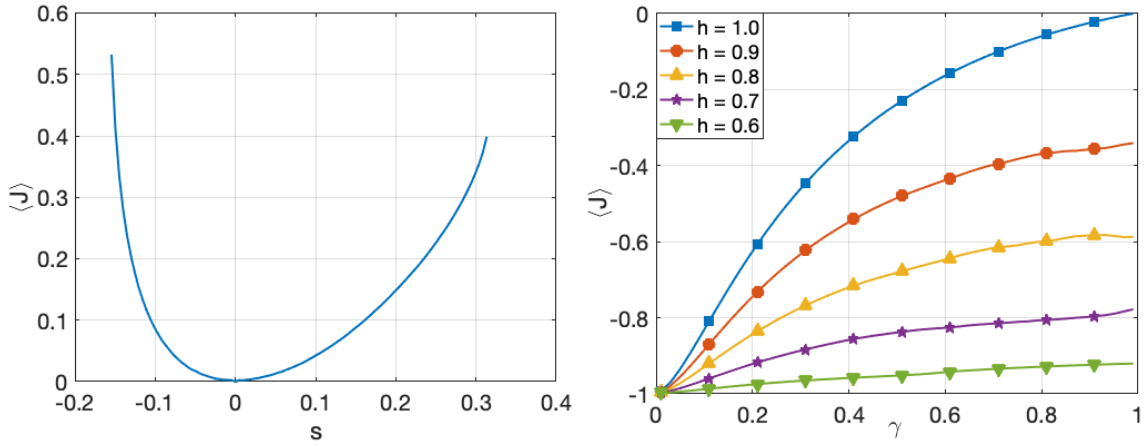

Figure 2-5: Long-time averaged behavior with respect to the map parameter for the sawtooth (left) and cusp map (right). The objective function itself does not depend to the parameter, and is defined as $J(x) = \cos(2\pi x)$. In our computations, J is averaged over 100 million samples.

systems are included in Chapter 3.

The above description can be mathematically rephrased as follows. For almost every x_0 uniformly distributed in $(0, 1)$,

$$\langle J \rangle = \lim_{N \to \infty} \frac{1}{N} \sum_{i=0}^{N-1} J(x_i) = \int_U J(x)\, \rho(x)\, dx. \tag{2.7}$$

Thus, in ergodic systems, there exist two alternative ways of computing the long-time average, either through the averaging of the time-series or ensemble averaging. However, the latter requires prior computation of the probability distribution, which is generally expensive. Given these preliminary concepts, we will now review the 1D version of Ruelle's formula, which we later use to compute the derivative of $\langle J \rangle$ with respect to the map parameter.

2.2 Sensitivity analysis of one-dimensional maps

Ruelle rigorously derived a formula for the derivative of the quantity of interest, $\langle J \rangle$, with respect to the controllable parameter s [148]. This expression is an ensemble average (or expectation) with respect to ρ, which can be simplified for one-dimensional

maps $\varphi : U \to U$ to

$$\frac{d}{ds}\int_U J(x)\,\rho(x)\,dx = \sum_{k=0}^{\infty}\int_U f(x)\,\frac{d\big(J\circ\varphi_k\big)(x)}{dx}\,\rho(x)\,dx, \tag{2.8}$$

where

$$f(x) := \frac{\partial\varphi\big(\varphi^{-1}(x)\big)}{\partial s} \tag{2.9}$$

reflects the parametric perturbation of the map, while U refers to the unit interval $[0, 1)$. In the subsequent chapters, we use χ to represent f in multi-dimensional systems.

A direct evaluation of Eq. 2.8 is computationally cumbersome for the following reason. Notice that the integrand of the right hand side involves a derivative of the composite function that can be expanded using the chain rule to the form

$$\frac{d(J\circ\varphi_k)(x)}{dx} = \Big(\frac{dJ}{dx}(\varphi_k(x))\Big)\prod_{j=0}^{k-1}\frac{\partial\varphi}{\partial x}(\varphi_j(x)). \tag{2.10}$$

As discussed in Chapter 1, the product of map derivatives exponentially grows with k, which means that

$$\frac{d\big(J\circ\varphi_k\big)}{dx}(x) \sim \mathcal{O}(e^{\lambda k}), \tag{2.11}$$

Nevertheless, Ruelle's series still convergences, as large quantities cancel out along and/or across infinitely long trajectories. Generating several trajectories and truncating them at a sufficiently small k is also hopeless, because the mean-squared error decreases proportionally to time if $\mathcal{O}(e^{2\lambda k})$ trajectories are available [29]. Thus, the direct evaluation of Ruelle's formula is computationally infeasible. For example, consider the sawtooth map with $s \in [-\frac{1}{2\pi}, \frac{1}{\pi}]$. In this case, $(\partial\varphi/\partial x) \in [1, 4]$. One can easily verify that even for moderate values of k, an overflow error is quickly encountered. Another challenge is that the evaluation of the SRB distribution requires expensive computation of map probability densities [19]. In recent studies [32, 31], Ruelle's formula has been reformulated to a different ensemble average, known as the S3 formula. The latter has been derived for maps with one positive LE, and is based

on splitting of the total sensitivity into stable and unstable contributions. Note the notion of splitting of the perturbation space is irrelevant for 1D maps, and the one-dimensional perturbation is, by the definition of chaos, unstable. Therefore we will skip some aspects of the original derivation, and note that our derivation represents only the unstable component of sensitivity in [32] specialized to 1D.

Motivated by [32, 31], we regularize Ruelle's formula defined by Eq. 2.8–2.9 by eliminating the problematic differentiation operator from the composite function. The new sensitivity expression has the following form,

$$\frac{d}{ds}\int_U J(x)\,\rho(x)\;dx = -\sum_{k=0}^{\infty}\int_U \nabla_\rho f(x)\,J\big(x_k\big)\;\rho(x)\;dx, \tag{2.12}$$

where

$$\nabla_\rho f(x) := \frac{1}{\rho(x)}\frac{d\big(\rho(x)\,f(x)\big)}{dx} = \frac{df}{dx}(x) + f(x)\,g(x), \tag{2.13}$$

and,

$$g(x) := \frac{1}{\rho(x)}\,\frac{d\rho}{dx}(x). \tag{2.14}$$

For one-dimensional maps, the derivation is simple, as it requires integrating Eq. 2.8 by parts and the fact that the integral of df/dx at the boundary of U vanishes; see Appendix 2.6.2 for the full derivation. We observe that both J and df/dx have their analytical forms. However, the function $g(x)$, which will be referred to as *density gradient* or *SRB measure gradient*, does not have a closed-form expression, since the SRB distribution ρ is unknown. The density gradient g represents the variation in phase space of the logarithm of $\rho(x)$,

$$g(x) := \frac{1}{\rho(x)}\,\frac{d\rho}{dx}(x) = \frac{d\log\rho(x)}{dx}. \tag{2.15}$$

In the next section, we focus on further interpretations of $g(x)$, its computation and verification on the 1D maps introduced in Section 2.1.

2.3 Computing SRB density gradients of one-dimensional maps

In this section, we focus on the density gradient function, denoted by $g(x)$. First, we present a computable, iterative scheme for $g(x)$. Moreover, we provide an intuitive explanation for $g(x)$ and visualize it on the maps introduced in Section 2.1.

The $g(x)$ function satisfies the following recursive relation

$$g\big(\varphi(x)\big) = \frac{g(x)}{d\varphi(x)/dx} - \frac{d^2\varphi(x)/dx^2}{\big(d\varphi(x)/dx\big)^2}, \tag{2.16}$$

which we derive in Appendix 2.6.3. This recursive procedure can be used to approximate $g(x)$ along a trajectory in the asymptotic sense. Indeed, we need a sufficiently large number of iterations to obtain an accurate approximation of $g(x)$ [162]. In practice, we generate a sufficiently long trajectory, compute first and second derivatives of the map evaluated along that orbit, and apply Eq. 2.16. We arbitrarily set $g(x_{\text{init}}) = 0$, to start the recursive procedure, and obtain $g(\varphi(x_{\text{init}}))$. The recursion is continued by setting $x = \varphi(x_{\text{init}})$, and so on. For a sufficiently large K, the true value of $g(\varphi_K(x_{\text{init}}))$ is approached, for almost every initial condition x_{init}. In Chapter 5, we prove that the recursion for g converges in uniformly hyperbolic systems.

2.3.1 Interpretation of the density gradient iterative formula

To intuitively understand the density gradient formula in Eq. 2.16, we isolate the effect of each term. Let us consider a small preimage subset of φ and two different mapping types: 1.) linear (straight line) and 2.) constant curvature mapping (e.g., quadratic function). These two cases are graphically shown on the left- (numbered as 1) and right-hand (numbered as 2) side of Figure 2-6, respectively. The x-axis represents an interval around x_k (chosen preimage subset), while the y-axis an interval around $x_{k+1} = \varphi(x_k)$ (corresponding image). The density ρ, around each interval, is shown adjacent to the axes, as a colormap. Specific colors reflect the distribution of ρ on a logarithmic scale.

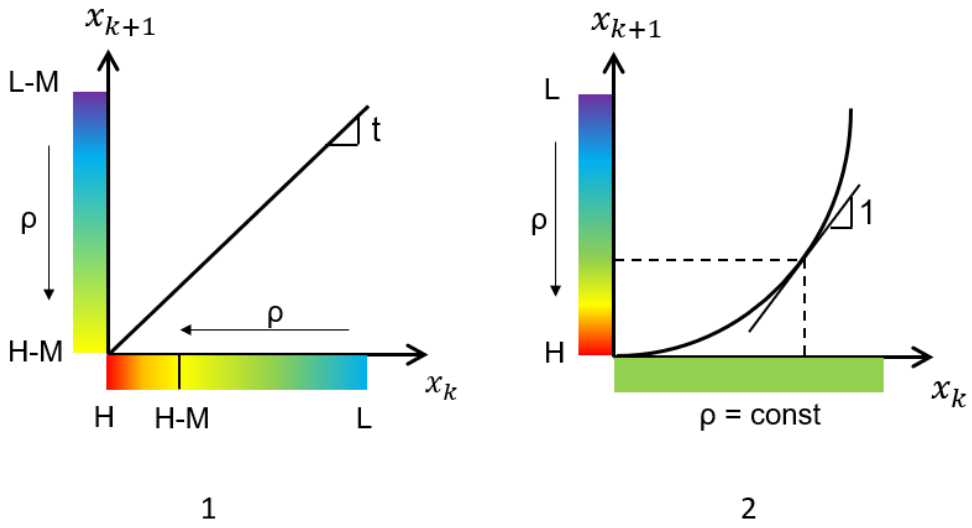

Figure 2-6: Graphical representation of two different scenarios in one-dimensional maps, to intuitively understand the derivation of g. The bold lines illustrate the map, while shaded bars adjacent to each axis represent the corresponding density distributions on a logarithmic scale. The region around H corresponds to a high value of density, while the region around L to low values. The slope of the line is indicated as t.

1. Let $(x-\epsilon, x+\epsilon)$ be a subset of the system's preimage at kth time step. Under the assumption of linear transformation, the first derivative of φ is constant within that subset. We additionally assume that the Jacobian is positive, while the density decreases monotonically; thus, $\rho(x+\epsilon) < \rho(x-\epsilon)$. Due to the property of measure invariance (probability mass conservation), the following relation[1]:

$$\rho(\varphi(x)) = \frac{\rho(x)}{|d\varphi/dx|} \tag{2.17}$$

holds. This equation implies that the densities at the neighborhood of given point on the attractor and its preimage are simply related through a constant number. Notice that

$$\log \rho(\varphi(x)) = \log \rho(x) - \log \frac{d\varphi}{dx}. \tag{2.18}$$

[1]This is a version of the density evolution equation involing the Frobenius-Perron operator. See Chapter 4 for more details.

Thus, on a logarithmic scale, the density distribution of the preimage is shifted by a constant compared to the current state. The colorbars of Figure 2-6 (1) illustrate this behavior. The initial density range, bounded by L and H, is uniformly shifted to the left by the value of M. The crux of this study case is the following observation. Since $d\varphi/dx$ is constant, rates of change of logarithmic densities in their respective coordinate systems, x_k and $x_{k+1} = \varphi(x_k)$, are the same. In addition, the preimage interval is expanded by a factor of $d\varphi/dx$. A mathematical equivalent of the two preceding statements can be obtained by differentiating Eq. 2.18,

$$\left(\frac{1}{\rho}\frac{d\rho}{dx}\right)\Bigg|_{\varphi(x)}\frac{d\varphi}{dx} = \left(\frac{1}{\rho}\frac{d\rho}{dx}\right)\Bigg|_{x}. \tag{2.19}$$

From the definition of g, this reduces to

$$g(\varphi(x)) = \frac{g(x)}{d\varphi/dx}, \tag{2.20}$$

which is consistent with Eq. 2.16 by setting $d^2\varphi/dx^2 = 0$.

2. In the second scenario, we aim to isolate the effect of the second term of Eq. 2.16. Thus, we will consider uniform measure distribution at the preimage stage, i.e., at step k. Thus, $g(x_k)$ is zero by definition, which is not necessarily true for $g(x_{k+1})$. We also assume that φ is a monotonically increasing quadratic function intersecting the origin. Due to the probability mass conservation (see Eq. 2.18), ρ is re-scaled by $d\varphi/dx$ after one iteration of φ. Now, let us specifically focus on the preimage point x at which the slope is exactly one, as represented graphically in Figure 2-6. We observe that at the image of the region to the left of x the density increases and *vice versa*. Indeed, the larger the first derivative of the map, the lower the density after transformation. This effect can be mathematically described by taking the derivative of Eq. 2.18 with respect to

x and applying the definition of g,

$$g(\varphi(x)) = -\frac{d^2\varphi/dx^2}{(d\varphi/dx)^2}. \tag{2.21}$$

In conclusion, the evolution of measure gradient is not only affected by expansion/compression factor, represented by $d\varphi/dx$, but it also takes into account the rate of volume change, represented by $d^2\varphi/dx^2$. Both effects directly follow from the measure preservation or probability mass conservation laws. In Chapter 3, these concepts are generalized and explained in more detail.

2.3.2 Numerical examples of density gradients

In the second part of this section, we show numerical results of the density gradient procedure in two examples, the sawtooth and cusp maps, which were introduced in Eq. 2.3 and Eq. 2.5 respectively. Figure 2-7 shows the stationary probability densities of the sawtooth map at different values of s. We observe that all curves appear differentiable, however their derivatives are large, near the interval boundaries, when s is close to $-1/(2\pi)$ or $1/\pi$. In Figure 2-8 we show the distribution of the (averaged) density gradient function, $g(x)$, computed using Eq. 2.16, at different values of s, and compare it against its finite difference approximation: $(\log(\rho(x+\epsilon)) - \log(\rho(x-\epsilon)))/(2\epsilon)$.

Note that the expected value of the density gradient is always zero since

$$\int_U g(x)\ \rho(x)\ dx = \int_U \frac{\partial \rho}{\partial x} dx = [\rho(x)]_0^1 = 0. \tag{2.22}$$

We repeat a similar experiment for the cusp map, whose results are presented in Figures 2-9–2-10. We observe a behavior similar to the sawtooth map.

Similarly to the sawtooth density, the densities computed for the cusp map appear to be differentiable over a range of the parameter γ. However, as γ approaches the value of 1, ρ acquires large slopes at the boundaries. The boundedness of $d\rho/dx$ is needed for the computation of $g(x)$ to be well-conditioned.

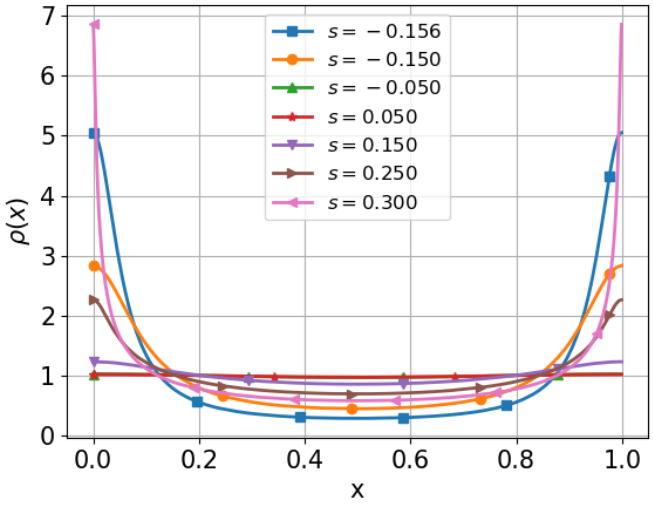

Figure 2-7: The plot shows the empirically estimated stationary probability distributions achieved by the sawtooth map (Eq. 2.3). Every curve was generated using 125,829,120,000 samples and counting the number of solutions in each of 2048 bins of equal length in the interval $[0, 1)$.

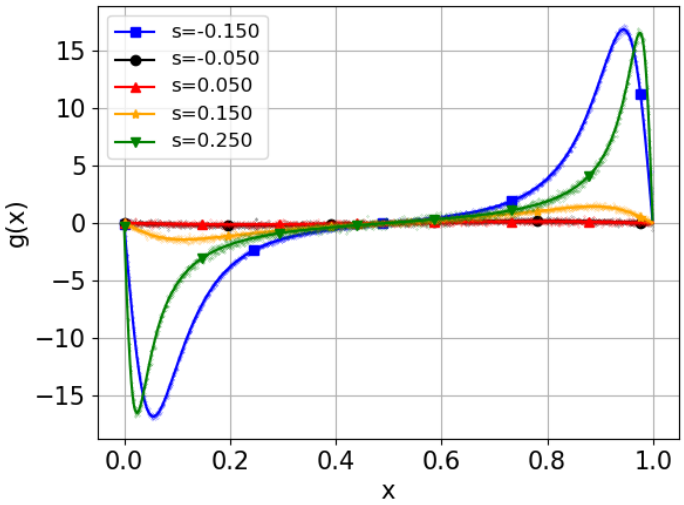

Figure 2-8: Density gradient function, $g(x)$ (solid lines), generated using Eq. 2.16 and compared against the empirically computed value of $g(x)$ (dots), where the derivative of $\rho(x)$ is estimated using finite difference.

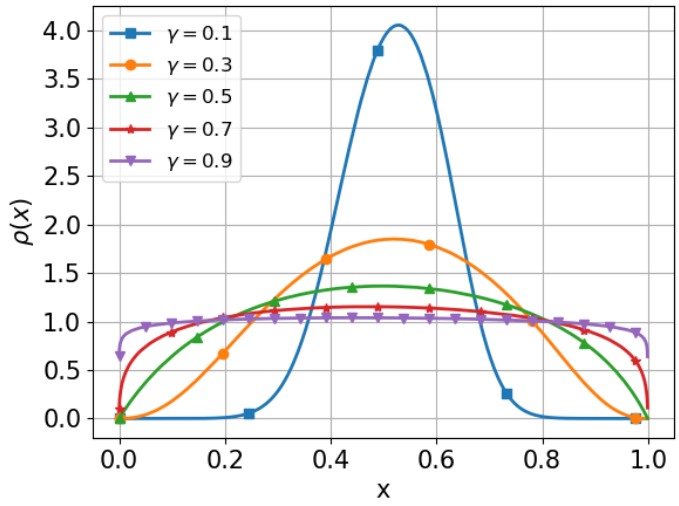

Figure 2-9: The plot shows the empirically estimated stationary probability distributions achieved by the cusp map (Eq. 2.5), at $h = 1$ and the indicated value of γ. All curves were generated in the same fashion as for the sawtooth case.

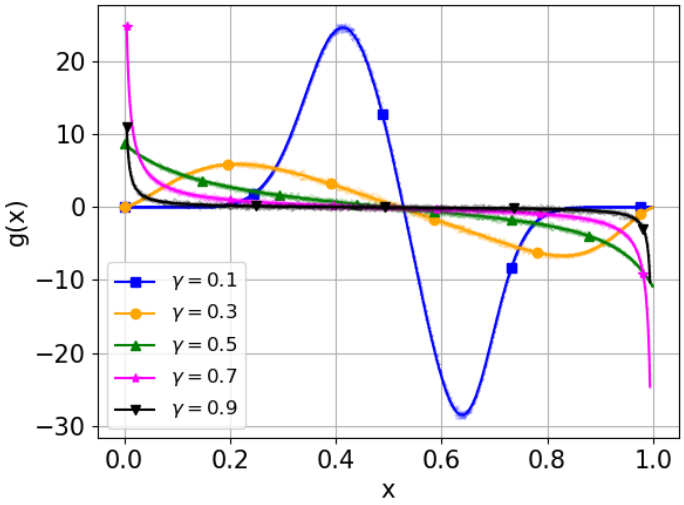

Figure 2-10: The plot compares $g(x)$ (solid lines) against the derivative of the empirically estimated stationary probability distributions (dots) achieved by the cusp map (Eq. 2.5), at $h = 1$ and indicated value of γ. All curves were generated in the same fashion as for the sawtooth case.

2.4 Sensitivity as a sum of time correlations

The main focus of this chapter is the evaluation of Eq. (2.12), which is an infinite series of k-lag time correlations between $\nabla_\rho f$ and J. Here, the k-lag time correlation is defined as the ensemble average of the product of two quantities with respect to ρ, where one is evaluated k time steps into the future. Due to the assumed ergodic behavior of φ, these ensemble averages can be approximated by computing time averages of data collected along truncated orbits initiated at a random initial condition. Therefore, the sought-after sensitivity can be approximated as follows,

$$\frac{d}{ds}\int_U J(x)\,\rho(x)\,dx \approx -\frac{1}{N}\sum_{k=0}^{\infty}\sum_{n=0}^{N-1}\Big(\frac{df}{dx}(x_n) + f(x_n)g(x_n)\Big)\,J\big(x_{n+k}\big), \tag{2.23}$$

where $x_{n+1} = \varphi(x_n)$, and g is generated "in parallel" using the recursion from Eq. 2.16. Systems in which the lag-k time correlations converge to zero as $k \to \infty$ are known as *mixing systems* [8]. In Axiom A mixing systems [164], the rate of decay of time correlations is exponential [194]. One can prove that one-dimensional maps are Axiom A systems if their derivatives are different than 1 everywhere on the attractor. The family of sawtooth and cusp maps that we described satisfy this requirement, which means that we need to compute a relatively small number of time correlations to obtain accurate approximations of the sensitivity. We also alert that one can empirically observe the exponential decay of correlations only if the generated trajectories are sufficiently long (see Chapter 5) for a detailed analysis.

2.4.1 Examples of sensitivity computation

To numerically verify Eq. 2.23, we consider a family of indicator functions, which will serve as observables $J_c(x)$ (objective functions). Each of them will be denoted as $J_c(x) := \delta_c(x)$ and defined such that its value is a constant 1 in a small interval around c and zero everywhere else on the unit interval. Note that every Riemann-integrable function can be approximated, to arbitrary precision, by linear combinations of such indicator functions. With this particular choice, Eq. 2.23 gives us the gradient of the

probability density, since

$$\frac{d}{ds}\int_U \delta_c(x)\,\rho(x;s)\,dx = \int_U \delta_c(x)\frac{\partial\rho(x;s)}{\partial s}\,dx \approx \frac{\partial\rho(c;s)}{\partial s}. \tag{2.24}$$

By running our linear response algorithm at several values of c in the interval $[0,1)$, we effectively obtain an approximation of $d\rho/ds$ over U. The reason we pick indicator functions as observables is that we can easily compare the result of our new scheme against the finite difference approximation of $d\rho/ds$. Since we deal with one-dimensional maps, the latter can be computed by approximating the density empirically at two different values of s and applying, for example, the central finite difference scheme. These estimates can then be used to calculate the derivatives of any other function with a differentiable statistics, using numerical quadrature. The choice of indicator functions exhibits yet another advantage of the regularized method over Ruelle's formula (Eq. 2.8). Indeed, the former is also applicable to objective functions that have non-differentiable points, since unlike in a direct evaluation of Ruelle's formula, the derivative of $J(x)$ is never used.

Figure 2-11 shows numerical results for the cusp map, in which the linear response is computed using the new formula (Eq. 2.23) and compared with the central difference derivative. We observe that only a few terms of the series are required to produce accurate sensitivities. Figure 2-12 clearly indicates that the consecutive terms of the series in Eq. 2.23 exponentially decay in norm. We repeat a similar experiment for the sawtooth map (see Figures 2-13 and 2-14). In this case, we only need three terms of Eq. 2.23 to obtain a result that is indistinguishable from its finite difference approximation. The consecutive terms of Eq. 2.23 also decay exponentially in norm, which is consistent with the theory.

2.4.2 Computational performance of the linear response algorithm

Finally, we compare the linear response method and classical finite difference method in terms of computational efficiency. Figures 2-15–2-16, generated respectively for

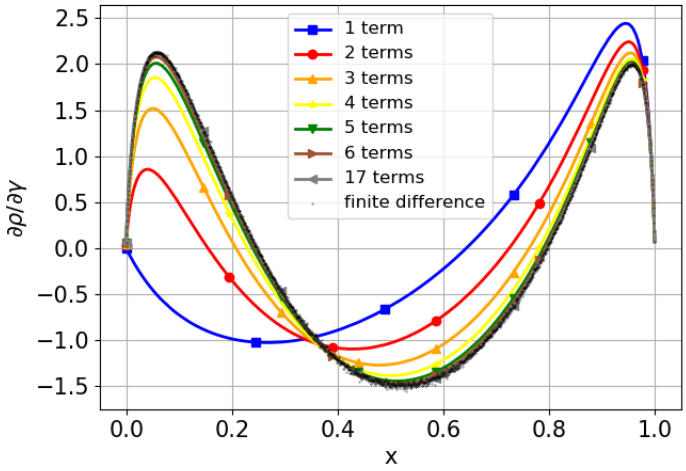

Figure 2-11: Sensitivity of the density of the cusp map with respect to γ at $h = 1, \gamma = 0.5$. The solid lines represent the result of Equation (2.12) when a finite number of terms is used in the summation over k. The solid line marked with ($\triangleleft$) represents Equation (2.12) evaluated with 17 terms, which is visibly indistinguishable from the same series summed over 6 or more terms. The dots represent the finite difference derivative of the density, evaluated based on the empirical density at $h = 1, \gamma = 0.505$ and at $h = 1, \gamma = 0.495$. Each quantity is evaluated with 125,829,120,000 samples.

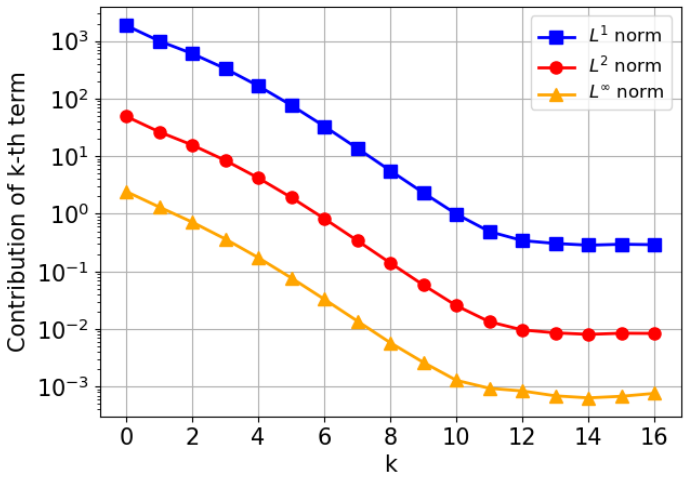

Figure 2-12: Contributions from the k-th term to Equation (2.12) for the cusp map. Later terms are overwhelmed by statistical noise.

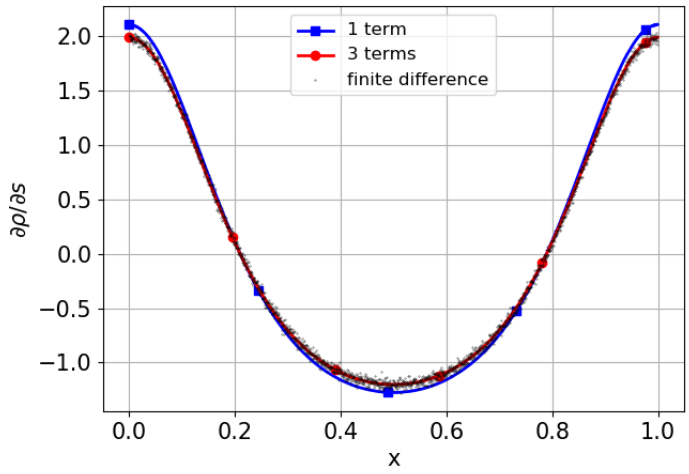

Figure 2-13: Sensitivity of the density of the sawtooth map with respect to s at $s = 0.1$. The solid lines represent the result of Eq. 2.12 when a finite number of terms is used in the summation over k. The solid line marked with (◯) represent Equation (2.12) evaluated with 3 terms, is aligned with the the finite difference derivative of the density, evaluated based on the density at $s = 0.105$ and at $s = 0.095$. Each quantity is evaluated with 125,829,120,000 samples.

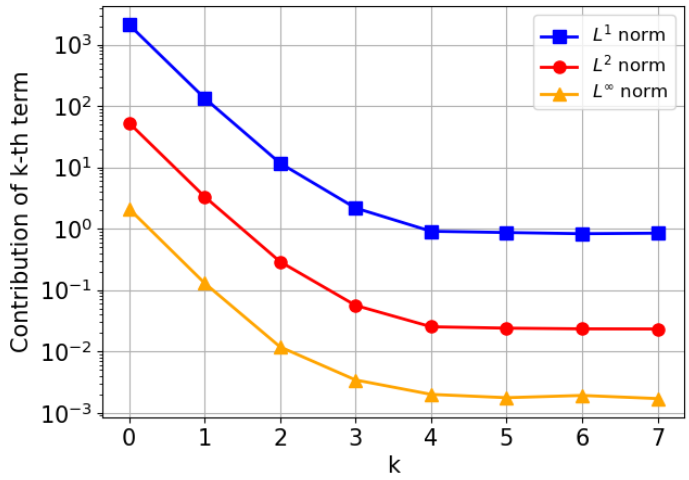

Figure 2-14: Contributions from the k-th term to Eq. 2.12, for the sawtooth map. Later terms are overwhelmed by statistical noise.

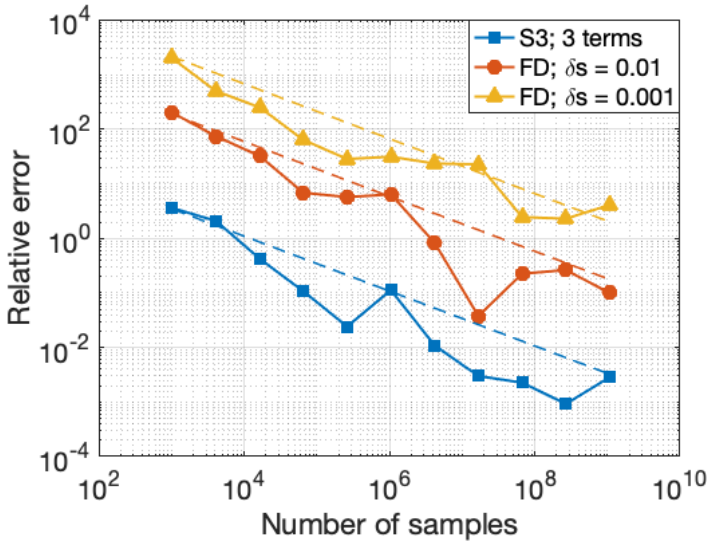

Figure 2-15: Relative error of the linear response and finite difference methods as a function of the trajectory length. We compute the parametric derivative of density of the sawtooth map at $s = 0.1$ on the left boundary ($x = 0$). In the linear response computation (curve marked with ($\square$)), we consider only first three terms of Eq. 2.12, which corresponds to the line marked with ($\bigcirc$) in Figure 2-13. For the finite difference approximation, we calculate densities at $s = 0.105, 0.095$ (curve marked with ($\bigcirc$)) and $s = 0.1005, 0.0995$ (curve marked with ($\triangle$)). In addition, we also computed the linear response with our algorithm using 125,829,120,000 samples and 9 terms of Eq. 2.12, which serves as a reference value. The dashed lines are proportional to the inverse of the square root of the number of samples.

the sawtooth and cusp maps, clearly indicate that the proposed method outperforms its competitor. We observe that the former requires a few orders of magnitude fewer samples to guarantee a similar error. This is a very promising observation in the context of analysing higher-dimensional systems, since the large cost of generating very long trajectories could make such computations infeasible.

Note in the case of both the linear response and finite difference methods, the error is upper-bounded as follows [32],

$$\text{error} \leq \frac{C}{\sqrt{N}}, \tag{2.25}$$

where N denotes the number of samples, while C is some positive number. This means we observe a convergence rate of a typical Monte Carlo simulation in both

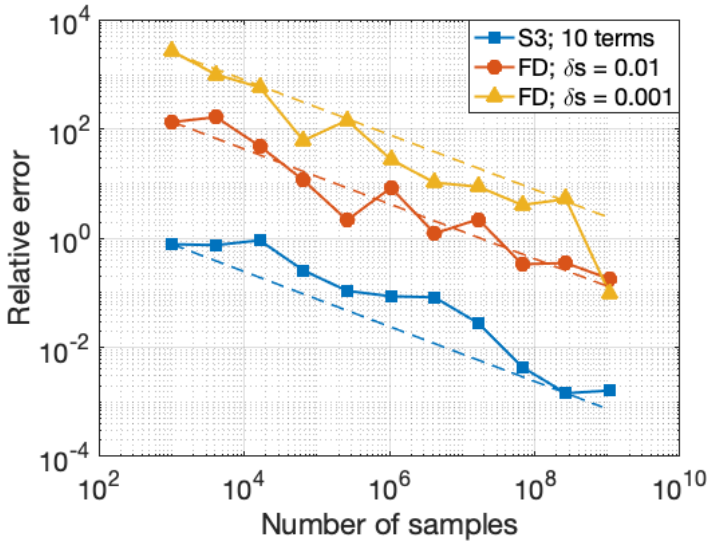

Figure 2-16: Relative error of the linear response and finite difference methods as a function of the trajectory length. We compute the parametric derivative of density of the cusp map at $h = 1$, $\gamma = 0.5$ in the middle of the domain U ($x = 0.5$). In the linear response computation (curve marked with ($\square$)), we consider first ten terms of Eq. 2.12. For the finite difference approximation, we calculate densities at $\gamma = 0.505, 0.495$ (curve marked with ($\bigcirc$)) and $\gamma = 0.5005, 0.4995$ (curve marked with ($\triangle$)). In addition, we also computed the linear response with our algorithm using 125,829,120,000 samples and 17 terms of Eq. 2.12, which serves as a reference value. The dashed lines are proportional to the inverse of the square root of the number of samples.

methods. However, the factor C is substantially larger in case of finite differencing. Moreover decreasing the step size (indicated as δs) in the finite difference calculation, worsens the accuracy, due the dominance of statistical noise.

2.5 Summary

We demonstrate a novel method based on Ruelle's formalism to compute the linear response, i.e., the parametric derivative of long-time statistics, in one-dimensional chaotic systems. To eliminate the problem of *butterfly effect*, the new method regularizes Ruelle's expression through partial integration. One of the by-products of integration by parts is a quantity called *density gradient*, defined as the derivative of the logarithmic density with respect to the state. This quantity plays a key role

in the computation of the linear response, and is efficiently computed in a recursive manner thanks to the property of ergodic measure conservation.

The proposed linear response method requires evaluating an infinite series of time correlations between regularized quantities. In uniformly hyperbolic systems, the time correlations are guaranteed to decay exponentially fast as the time gap increases. It means that a relatively small number of the series terms should be computed to accurately estimate the sensitivity, which we verify in our numerical tests. We also demonstrate that the linear response formula requires several orders of magnitude less samples compared to the finite difference method. This benefit stems precisely from the regularization of Ruelle's original expression.

While this work is restricted to simple one-dimensional expansive maps, the presented examples illustrate the concepts that will be later used in the derivation of the generalized linear response algorithm, known as the *space-split sensitivty* (S3) method. Here we also point out that, in certain cases, the slopes of ergodic densities could be very large locally, which corresponds to heavy tailedness of the density gradient distribution. This phenomenon, as well as systematic methods for differentiating SRB densities on higher-dimensional unstable manifolds is the main topic of the next two chapters.

2.6 Appendices

This section includes supplementary material supporting certain arguments presented in this chapter.

2.6.1 Binary floating point problem in simulating 1D maps

Consider the case $s = 0$. Map 2.3 can be compactly expressed using the modulo operator, i.e. $x_{n+1} = 2x_n \bmod 1$. It means we multiply x_n by 2 and if $x_{n+1} > 1$, then we also subtract 1. Using floating point arithmetic, we will observe that there exist $N > 0$ such that $x_n = 0$ for all $n \geq N$, which contradicts the assumption of chaotic behavior. This phenomenon is due to the round-off errors associated with the

modulo operator. To circumvent this problem, one can change the divisor parameter (of the modulo operation) from 1 to $1 - \epsilon$, where ϵ is a small number, e.g. $\epsilon = 10^{-6}$. Another possible (and simple) workaround might be a change of variables such that the domain of the new variable has irrational length. Note this approach would also require a modification of the objective function.

2.6.2 Derivation of the linear response formula for 1D maps

In this section, we will show Eq. 2.8–2.9 are equivalent to Eq. 2.12–2.14. Throughout this derivation we will use a short-hand notation for the composition $v \circ \varphi_k = v_k$, where v is some scalar function defined on $U = (0, 1)$, while k is some integer. If $k = 0$, the subscript is dropped. First, note

$$\int_U f \; \frac{dJ_k}{dx} \; \rho \; dx = \int_U \frac{d}{dx} \left(f \; J_k \right) \rho \; dx - \int_U J_k \frac{df}{dx} \; \rho \; dx. \tag{2.26}$$

Integrate the first term of Eq. 2.26 by parts,

$$\int_U \frac{d}{dx} \left(f \; J_k \right) \; \rho \; dx = \left[f \; J_k \; \rho \right]_{U_L}^{U_R} - \int_U f \; J_k \frac{\partial \rho}{\partial x} \; dx, \tag{2.27}$$

where $U_L = 0$ and $U_R = 1$ correspond to the left and right boundary of U, respectively. Since the domain is periodic, the first term of Eq. 2.27 vanishes. Thus, we can combine Eq. 2.26 and Eq. 2.27 to conclude that

$$\int_U f \; \frac{dJ_k}{dx} \; \rho \; dx = - \int_U J_k \left(\frac{\partial f}{\partial x} + \frac{1}{\rho} \frac{\partial \rho}{\partial x} \right) \; \rho \; dx. \tag{2.28}$$

2.6.3 Derivation of an iterative procedure for the ergodic measure gradient in 1D maps

The purpose of this section is to derive the iterative procedure to calculate the density gradient g. We use the same notational convention as in Appendix 2.6.2. Let us

consider a function h that is integrable in $U = (0, 1)$ and vanishes at $U_L = 0$ and $U_R = 1$. Using the definition $g = (1/\rho)(\partial\rho/\partial x)$, and integrating by parts, we obtain

$$\int_U g\ h\ \rho\ dx = \int_U h\ \frac{d\rho}{dx}\ dx = [h\ \rho]_{U_L}^{U_R} - \int_U \frac{dh}{dx}\ \rho\ dx = -\int_U \frac{dh}{dx}\ \rho\ dx. \tag{2.29}$$

The key property used in this derivation is the density preservation of φ. The map φ is density-preserving with respect to the density ρ if

$$\int_U f\ \rho\ dx = \int_U f \circ \varphi_k\ \rho\ dx$$

holds for any integer k and any scalar observable f. This implies the left hand side of Eq. 2.29 can be expressed as

$$\int_U g\ h\ \rho\ dx = \int_U g_1\ h_1\ \rho\ dx. \tag{2.30}$$

We now apply the density preservation together with the chain rule to the right hand side of Eq. 2.29, which gives rise to

$$-\int_U \frac{dh}{dx}\ \rho\ dx = -\int_U \left(\frac{dh}{dx}\right)_1\ \rho\ dx = -\int_U \frac{dh_1}{dx}\ \frac{1}{d\varphi/dx}\ \rho\ dx. \tag{2.31}$$

Note

$$\frac{dh_1}{dx}\ \frac{1}{d\varphi/dx} = \frac{d}{dx}\left(\frac{h_1}{d\varphi/dx}\right) - h_1\frac{d}{dx}\left(\frac{1}{d\varphi/dx}\right) = \frac{d}{dx}\left(\frac{h_1}{d\varphi/dx}\right) + h_1\frac{d^2\varphi/dx^2}{(d\varphi/dx)^2}, \tag{2.32}$$

and, using $h_1(U_L) = h_1(U_R) = 0$, integrate by parts to get,

$$-\int_U \frac{d}{dx}\left(\frac{h_1}{d\varphi/dx}\right)\ \rho\ dx = -\left[\frac{h_1}{d\varphi/dx}\ \rho\right]_{U_L}^{U_R} + \int_U \frac{h_1}{d\varphi/dx}\frac{d\rho}{dx}\ dx = \int_U \frac{h_1}{d\varphi/dx}\frac{d\rho}{dx}\ dx. \tag{2.33}$$

Combine Eq. 2.31–2.33 to observe that

$$-\int_U \frac{dh}{dx}\ \rho\ dx = \int_U h_1\ \left(\frac{g}{d\varphi/dx} - \frac{d^2\varphi/dx^2}{(d\varphi/dx)^2}\right)\ \rho\ dx. \tag{2.34}$$

Finally, by combining Eq. 2.29,2.30, and 2.34, we obtain the following identity,

$$\int_U h_1 \; g_1 \; \rho \; dx = \int_U h_1 \; \left(\frac{g}{d\varphi/dx} - \frac{d^2\varphi/dx^2}{(d\varphi/dx)^2} \right) \; \rho \; dx, \tag{2.35}$$

from which we infer that

$$g_1 = \frac{g}{d\varphi/dx} - \frac{d^2\varphi/dx^2}{(d\varphi/dx)^2}. \tag{2.36}$$

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

Differentiating ergodic measures on unstable manifolds

Ergodic or Sinai-Ruelle-Bowen (SRB) measures are limiting stationary distributions describing the statistical behavior of chaotic dynamical systems. Directional derivatives of SRB measure densities conditioned on unstable manifolds are critical in the sensitivity analysis of hyperbolic chaos. These derivatives, known as the *SRB density gradients*, are by-products of the regularization of Lebesgue integrals appearing in the original *linear response* expression. In this chapter, we propose a novel trajectory-driven algorithm for computing the SRB density gradient defined for systems with high-dimensional unstable manifolds. We apply the concept of measure preservation together with the chain rule on smooth manifolds. Due to the recursive one-step nature of our derivations, the proposed procedure is memory-efficient and can be naturally integrated with existing Monte Carlo schemes widely used in computational chaotic dynamics. We numerically show the exponential convergence of our scheme, analyze the computational cost, and present its use in the context of Monte Carlo integration.

This chapter consists of two major parts. The first one, spanning Sections 3.1 – 3.5, introduces and derives mathematical machinery for differentiating densities implied by smooth coordinate charts describing generic manifolds. The key idea proposed in that part is to parameterize smooth manifolds using quantile functions. We also

point out that Lebesgue integration of derivatives of strongly-oscillatory functions is a recurring challenge in computational science and engineering. Integration by parts is an effective remedy for huge computational costs associated with Monte Carlo integration schemes. **The content of this part was published in *Applied Mathematics and Computation* by *Elsevier* [160].** The second part, spanning Sections 3.6 – 3.10, applies these concepts to effectively differentiate SRB measures on multi-dimensional unstable manifolds. **We acknowledge that the second part was first published in *SIAM Journal on Scientic Computing* in 2022 by the *Society for Industrial and Applied Mathematics (SIAM)* [162]. Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.**

3.1 Preliminaries

The fundamental theorem of calculus states that

$$\int_a^b \partial_x f(x)\, dx = f(b) - f(a),$$

where f is some smooth function defined on the compact interval $[a, b]$. This theorem is critical in many applications, including computational sciences [14, 172]. For example, if the function f is strongly oscillatory, a numerical quadrature on the left-hand side would require many points and much computation to obtain accurate results. Nevertheless, the fundamental theorem guarantees that the positive and negative derivatives of this oscillatory function largely cancel each other out. Indeed, one can simply compute the right-hand side directly, without truncation error.

As a generalization of the fundamental theorem, consider the integral of $\partial_x f$ over the same domain under some Lebesgue measure m, which is an antiderivative of the density function ρ (a.k.a. the Radon-Nikodym derivative [123]), i.e., $dm(x) = \rho(x)\, dx$. In the classical version of the theorem, as mentioned in the first paragraph, the density ρ is constant and equals $1/(b-a)$ everywhere on the domain. If this is not the case,

however, the integration by parts of $\partial_x f$ involves the derivative of ρ,

$$\int_a^b \partial_x f(x)\, dm(x) = f\rho\Big|_a^b - \int_a^b f(x)\, \partial_x \rho(x)\, dx = f\rho\Big|_a^b - \int_a^b f(x)\, \frac{\partial_x \rho}{\rho}(x)\, dm(x) \quad (3.1)$$

Note also that Eq. 3.1 can be viewed as an extension of the Generalized Stokes Theorem to integrals involving non-uniform measures. The second term from the RHS of Eq. 3.1 does not appear in the classical version of this theorem because $\partial_x \rho$ vanishes everywhere on the domain. The integral in Eq. 3.1 can be approximated using a Monte Carlo integration scheme if a set of realizations of x, $\{x^1, x^2, ..., x^N\}$, distributed according to m, is given. However, if f is a strongly-oscillatory function with large magnitude, the Monte Carlo method applied directly to the integral on the left-hand side (LHS) of Eq. 3.1 would require a large amount of data to obtain an approximation with a reasonably small error [134, 117]. Alternatively, one can consider the right-hand side (RHS) of the same equation, which requires the function f itself, not its derivative. Assuming the density ρ is a well-behaved function, the variance of the integrand on the RHS is significantly smaller and, therefore, remarkably less data is needed to obtain an accurate result. However, extra computational effort must be put to evaluate $\partial_x \rho/\rho = \partial_x \log \rho$. The computation of that function, which we denote by g and call *density gradient*[1], is the main focus of this chapter.

Lebesgue integrals involving functions with high fluctuations are critical in the field of sensitivity analysis of chaotic dynamical systems. Ruelle [148, 150] derived a closed-form expression, known as the *linear response* formula, for the parametric derivative of the mean of a quantity of interest J. The linear response formula includes Lebesgue integrals of directional derivatives of a strongly oscillatory J over the manifold of a chaotic system. A computable version of Ruelle's formula, known as the space-split sensitivity (S3), was originally obtained through the integration by parts of the original formulation [32] and also presented in Chapter 2. The S3 algorithm was successfully applied in various low-dimensional systems in the computation [158]

[1]While the function g appearing in Sections 3.1 – 3.5 is not the SRB (ergodic) measure gradient, we will stick to the same notation due to the natural equivalence. Here, g represents a derivative of the logarithmic density implied by smooth measure transformation.

and assessment of existence [159] of parametric derivatives of statistical quantities describing chaos. The crux of the computation of the regularized Ruelle's formula is the *SRB density gradient*, defined as a directional derivative of the logarithm of the SRB density [195, 41] along the unstable manifold. While an efficient numerical procedure for the approximation of the SRB density gradient specialized to systems with one-dimensional unstable manifolds is available [32, 158, 159, 33], we still lack a generalizable algorithm applicable to arbitrary higher-dimensional chaotic systems.

The main purpose of the first part of this chapter is to derive a general formula for the density gradient g, defined on a differentiable m-dimensional manifold M immersed in the Euclidean space $\mathbb{R}^n$, $m \leq n$. In our analysis, we parameterize M using the chart $x(\xi) : \mathbb{R}^m \to \mathbb{R}^n$. Here, the g function is an m-element vector, whose i-th component equals a directional derivative of $\log \rho$, in the direction of a unit vector s_i, i.e.,

$$g_i = \frac{\partial_{s_i} \rho}{\rho} = \frac{\nabla_x \rho \cdot s_i}{\rho}.$$

The scalar function ρ is the density implied by $x(\xi)$. Without loss of generality, we assume that the i-th directional derivative is computed along the isoparametric line in the direction of increasing i-th component of ξ. Analogously to Eq. 3.1, the Lebesgue integral of the directional derivative of J over M with measure m can be written using g_i,

$$\int_M \nabla_x J(x) \cdot s_i(x)\, dm(x) = -\int_M J(x)\, g_i(x)\, dm(x), \tag{3.2}$$

where J is assumed to vanish on the boundary of M. For the reasons indicated above, it is computationally efficient to apply the Monte Carlo method to the RHS of Eq. 3.2. Analogous integration by parts is required to regularize the linear response [32]. Thus, the derivation of a computable expression for the density gradient defined on higher-dimensional smooth manifolds is a milestone in constructing algorithms for differentiating SRB measures. In addition, an explicit formula for g might serve as a valuable tool in general numerical procedures involving integrals over geometrically complex domains.

3.2 Computing density gradients on one-dimensional manifolds

In this section, we focus on the computation of the density gradient g in the simplest topological setting. In particular, we consider general smooth one-dimensional manifolds immersed in the Euclidean space $\mathbb{R}^n$ which can be described using a single parameter $\xi \in [0, 1]$. They are all geometrically represented by curves. Let us pick one out of infinitely many such curves and denote it by $\mathcal{C}$. We assume there exists a one-to-one map $x(\xi) \in \mathcal{C} \subset \mathbb{R}^n$, which is at least twice differentiable with respect to ξ, i.e., $x(\xi) \in C^2([0, 1])$. In this case, the density gradient function is a scalar quantity defined as a directional derivative along $\mathcal{C}$ of logarithmic density, $g = \partial_s \log \rho$, where $\rho : \mathcal{C} \to [0, 1]$ is a density function implied by $x(\xi)$. If we think of ξ as a realization of the random variable uniformly distributed in $[0, 1]$, then $x(\xi)$ is in fact the inverse cumulative distribution function (inverse CDF, a.k.a. the quantile function). Intuitively, $x(\xi)$ tells us that 100ξ % of all points mapped from the uniformly distributed set are located on the curve segment between $x(0)$ and $x(\xi)$. On the other hand, the density function ρ indicates the density of points mapped on $\mathcal{C}$ per unit curve length. Therefore, ρ is counter-proportional to the magnitude of the first derivative of $x(\xi)$. The intuitive relation between the probability distribution function and nonlinear coordinate transformation has been utilized in applied statistics (inference problems); see for example [57, 64]. These studies build upon the invariance property (transformation law) of the density function [17], which is reflected by the relation between ρ and the derivative of $x(\xi)$.

In the following three subsections, we analytically derive the expression for g in terms of the inverse CDF $x(\xi)$ for simple line manifolds, $n = 1$ (Section 3.2.1), and general curves, $n \geq 1$ (Section 3.2.3), and demonstrate its importance in a numerical integration experiment (Section 3.2.2). We illustrate all relevant concepts using a certain $x(\xi)$ associated with the Van der Pol equation,

$$\frac{d^2u}{dt^2} = 2(1 - u^2)\frac{du}{dt} - u, \quad u(0) = -a, \quad \frac{du}{dt}(0) = 0, \quad a > 0, \tag{3.3}$$

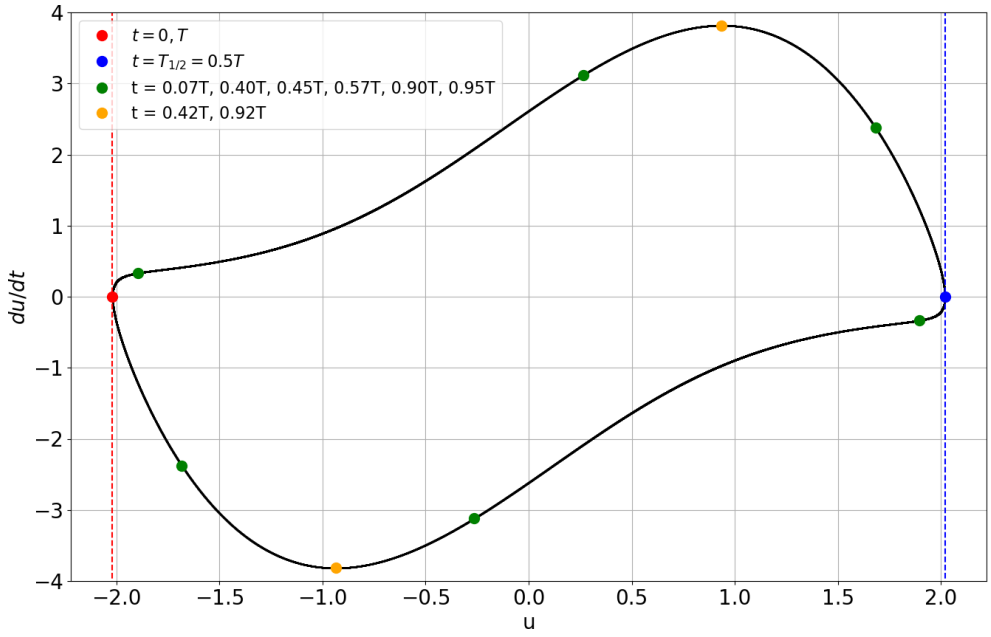

Figure 3-1: Trajectory of the Van der Pol oscillator (Eq. 3.3). The red dot represents the initial condition, as well as the solution after time T, while the blue dot indicates the solution after time $T_{1/2}$. The vertical dashed lines correspond to $u = -a$ and $u = a$ (boundaries of the range of u). At the green dots, the solution satisfies $du/dt + d^3u/dt^3 = 0$, while the zero acceleration state, $d^2u/dt^2 = 0$, is represented by orange dots. The solution to Eq. 3.3, represented by the solid black curve, is an example of a one-dimensional manifold immersed in $\mathbb{R}^2$.

which describes the coordinates of a 2D non-conservative oscillator with non-linear dumping [66]. In our numerical examples, we choose $a = 2.0199$, in which case the solution $[u(t), du(t)/dt]^T$ approximately lies on the limit cycle with period $T = 2T_{1/2} \approx 7.638$ and $u(t) \in [-a, a]$ for all $t \geq 0$. Figure 3-1 illustrates the limit cycle of Eq. 3.3, which has been computed using the second-order Runge-Kutta (midpoint) method with time step $\Delta t = 0.0001$. Indeed, the solution to the Van der Pol equation may serve as example of a nonlinear transformation from a parametric space (time t, in this case) to a smooth manifold immersed in $\mathbb{R}$ (in case of the $t \to u(t)$ mapping) or $\mathbb{R}^2$ (in case of the $t \to [u(t), du(t)/dt]^T$ mapping).

3.2.1 Line manifolds

We start from the simplest case, i.e., when $\mathcal{C}$ is a bounded line segment in $\mathbb{R}$, between a and b. The corresponding inverse CDF $x(\xi)$ differentiably maps $[0, 1]$ to $[a, b]$ and

is related to the density function by the following expression,

$$\xi(x) = \int_a^x \rho(y)\, dy \quad \forall x \in [a, b]. \tag{3.4}$$

Since $\xi \in [0, 1]$, $\rho(x)$ is in fact the probability density function (PDF) corresponding to the CDF $\xi(x)$, which satisfies $d\xi = \rho(x)\, dx$. Using the inverse function theorem, which asserts $f'(f^{-1}(c)) = 1/[(f^{-1})'(c)]$ for any differentiable one-to-one function f at any c such that $(f^{-1})'(c) \neq 0$, we conclude that

$$\frac{dx}{d\xi}(\xi)\, \rho(x(\xi)) = 1. \tag{3.5}$$

Eq. 3.5 indicates that at any point $x(\xi)$ on the manifold, the product of the PDF and derivative of the inverse CDF is constant. Thus, by differentiating Eq. 3.5 with respect to ξ and reshuffling terms, we obtain a direct expression for g at each point on the manifold,

$$g(x(\xi)) = \partial_x \log \rho(x(\xi)) = \frac{\partial_x \rho(x(\xi))}{\rho(x(\xi))} = -\frac{\frac{d^2x}{d\xi^2}(\xi)}{\left(\frac{dx}{d\xi}(\xi)\right)^2}. \tag{3.6}$$

To illustrate these functions and their relation, we will consider the solution to Eq. 3.3, $u(t)$, for $t \in [0, T_{1/2}]$, where $T_{1/2} \approx 3.819$. Based on Figure 3-1, it is evident that $u(t)$ is a one-to-one smooth function and $du/dt \geq 0$ in that time interval. In fact, we can apply a linear transformation $t \to \xi$ to notice that

$$x(\xi) = u\left(\xi T_{1/2}\right) \tag{3.7}$$

is a representation of the inverse CDF. Next, we compute the first and second derivative of Eq. 3.7 with respect to ξ and plug them into Eq. 3.6 to obtain the following formula for g along the trajectory,

$$g(u(t)) = -\frac{\frac{d^2u}{dt}(t)}{\left(\frac{du}{dt}(t)\right)^2} \overset{\text{Eq. 3.3}}{=} -\frac{2(1 - u^2(t))\frac{du}{dt}(t) - u(t)}{\left(\frac{du}{dt}(t)\right)^2}. \tag{3.8}$$

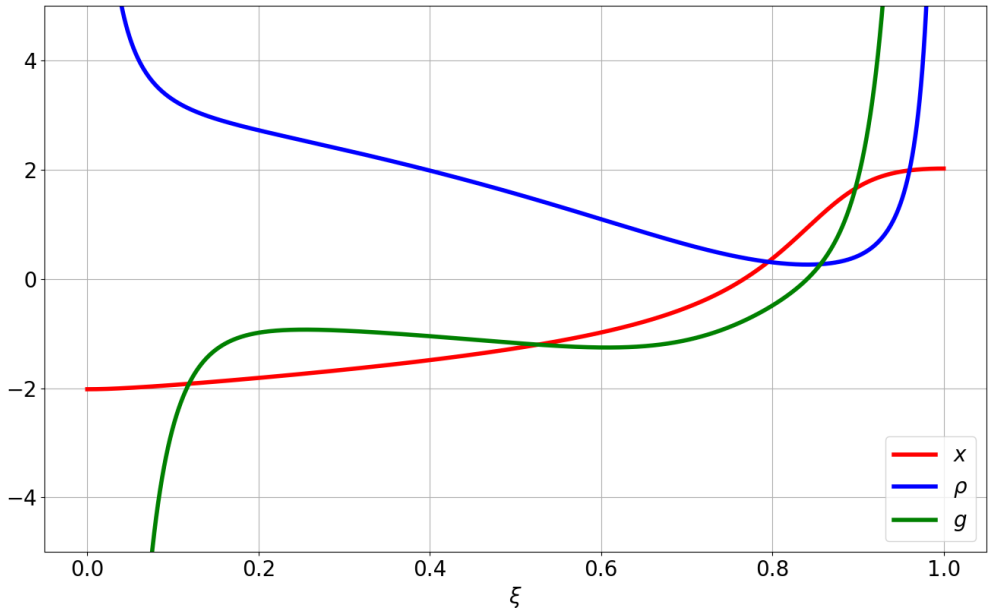

Figure 3-2: The inverse CDF function $x(\xi)$ defined by the solution to the Van der Pol equation, such that $x(\xi(t)) = u(t)$ for all $t \in [0, T_{1/2}]$ (red), and the corresponding density (blue) and density gradient function (green). We used data presented in Figure 3-1 to compute all the three functions.

We observe that the density gradient is invariant to any linear change of variables, i.e., when $d\xi/dt$ is constant. Given a numerical solution to Eq. 3.3, the density can be directly computed from

$$\rho(u(t)) = \frac{1}{(T_{1/2}\, du(t)/dt)},$$

which follows from Eq. 3.5, whereas the density gradient function can be evaluated using Eq. 3.8.

Figure 3-2 illustrates the inverse CDF $x(\xi)$, defined by Eq. 3.7, as well as the corresponding density and density gradient. We clearly observe that both ρ and g are undefined at the endpoints, i.e., at $\xi = 0$ and $\xi = 1$, which is a consequence of zero slope of $x(\xi)$. Moreover, the larger the rate of change of x, the smaller the value of ρ, which confirms our previous intuitive explanation of the density function. We also notice that the density gradient is zero at the point corresponding to a local extremum of ρ and the inflection point of $x(\xi)$.

3.2.2 Approximating integrals of a highly oscillatory function

We now demonstrate the use of the density gradient function in the numerical computation of a highly oscillatory function. Consider the following Lebesgue integral,

$$I = \int_{-a}^{a} \partial_x f(x)\, d\xi(x), \tag{3.9}$$

where $\xi(x)$ denotes a Lebesgue measure defined by Eq. 3.4, while f is a function whose first derivative is integrable and bounded. Certainly, it is assumed the above integral converges. Indeed, a sufficient condition for the convergence of I in this case is Lebesgue-integrability of the density gradient with respect to the density ρ [159], i.e., $g \in L^1(\rho)$. However, the necessary and sufficient condition imposes extra requirements for the f function itself, i.e., $\partial_x f \in L^1(\rho)$ or, equivalently, $\partial_x f\, \rho \in L^1([-a, a])$. In our experiment, the function f has the following form,

$$f(x) = \big((x-a)(x+a)\sin(Kx^2)\big)^2, \tag{3.10}$$

with some positive number K. We use Eq. 3.4 to rewrite the above integral, and then integrate it by parts. There exist a few scenarios when the resulting boundary term vanishes. One option is that the product $\partial_x f\, \rho$ is periodic and integrable on $[-a, a]$. Another possibility is when both $\partial_x f$ and ρ are bounded and at least one of them vanishes at the domain boundaries. In any case, two new versions of I, alternative to the original form (in Eq. 3.9), are available,

$$\int_{-a}^{a} \partial_x f(x)\, \rho(x)\, dx = I = -\int_{-a}^{a} f(x)\, g(x)\, d\xi(x). \tag{3.11}$$

To numerically approximate the integral I, we apply three distinct approaches. The integral in Eq. 3.9 and the RHS of Eq. 3.11 can be estimated using a Monte Carlo method, which requires generating a random sequence $\{x^1, x^2, ..., x^N\}$ distributed according to the measure ξ. If such a sequence is available, then the integral of any

Lebesgue-integrable function $h(x)$ can be approximated as follows,

$$\int_{-a}^{a} h(x)\, d\xi(x) \approx \frac{1}{N}\sum_{i=1}^{N} h(x^i), \tag{3.12}$$

since $\xi \in [0,1]$. Finally, the integral on the LHS of Eq. 3.11 is evaluated using a standard trapezoidal rule with a uniform N-element grid for x between $-a$ and a. In the calculation, we allot

$$x^i = u\left(\frac{i-1}{N-1}T_{1/2}\right). \tag{3.13}$$

It can be numerically verified that for this particular choice of the sequence, $g \notin L^1(\rho)$, but $\partial_x f \in L^1(\rho)$. It means that the integral I converges despite the blow-up of ρ and g at the boundaries of $[-a,a]$. To assess the Lebesgue-integrability of these functions, we applied the procedure described in Section 4 of [159]. This algorithm approximates the slope of the distribution tail of any function in the logarithmic scale.

In order to compare the performance of these three integration methods, we proceed as follows. First, we generate the sequence $\{x^1, x^2, ..., x^N\}$, $N = 10^5$ (time step is chosen such that $\Delta t = T_{1/2}/(N-1)$) and, using Eq. 3.5 and Eq. 3.8, we directly evaluate ρ and g at all points from that sequence. Subsequently, both the density and density gradient functions are linearly interpolated everywhere between $-a$ and a. We use these interpolators to approximate the two functions at any point of a uniform grid (trapezoidal rule) or sequence defined by Eq. 3.13 (Monte Carlo) for an arbitrary value of N. If K is sufficiently small, then the approximation error of the trapezoidal rule is expected to be upperbounded by $\mathcal{O}(1/N)$, because the integrand, $\partial_x f\, \rho$, is Lebesgue-integrable [43]. According to the Nyquist-Shannon sampling theorem, however, the discrete representation of the integrand may not be captured properly if K is very large, in which case the trapezoidal rule's error decays as in a typical Monte Carlo method. Figure 3-3 shows the behavior of the relative error of the approximation of I obtained using these three methods. The error is computed with respect to the reference solution obtained through the trapezoidal rule using $N = 10^8$ points.

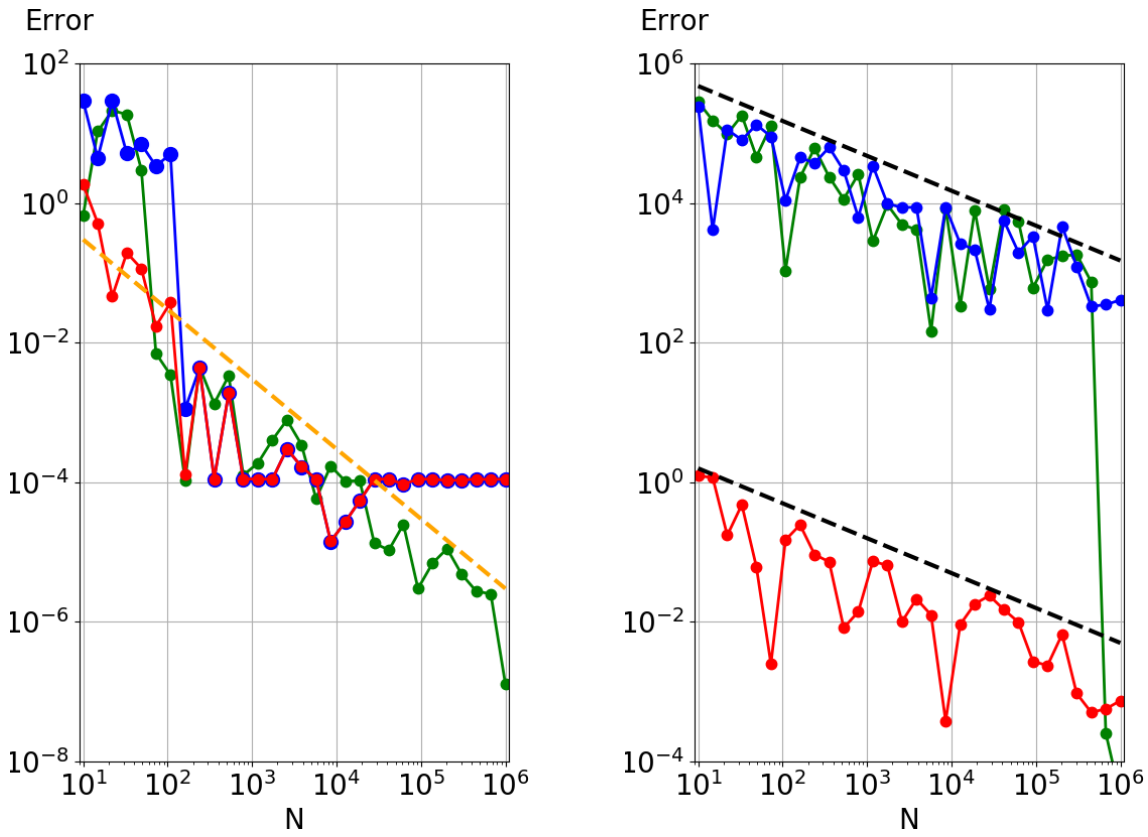

Figure 3-3: Relative error of the approximation of I for $K = 10$ (left) and $K = 100000$ (right) obtained using three methods: Monte Carlo integration applied to Eq. 3.9 (blue), Monte Carlo integration applied to the RHS of Eq. 3.11 (red), and trapezoidal rule applied to the LHS of Eq. 3.11 (green). Black and orange dashed lines are reference lines representing functions proportional to $N^{-1/2}$ and N^{-1}, respectively. In each of these plots, we computed the relative error with respect to the approximation of I obtained using the trapezoidal with $N = 10^8$ samples.

We observe that for a moderately-oscillatory integrand ($K = 10$), the relative error of the trapezoidal rule (green curve) decays as $\mathcal{O}(1/N)$, which confirms the theoretical estimates. In this case, the performance of both the Monte Carlo approximations (blue and red curves) does not differ much from the trapezoidal rule's. The Monte Carlo approximation clearly converges to a solution slightly different than the reference solution, which is a consequence of the fact the latter was generated using the trapezoidal rule for a linearly interpolated function. This example indicates that there is no reason to perform integration by parts and compute g to approximate integrals of low- or moderately-oscillatory functions. The right-hand side plot of Figure 3-3 corresponds to a different scenario, i.e., when f is highly-oscillatory ($K = 10^5$).

The error now decays $\mathcal{O}(1/\sqrt{N})$ at $N \in [10^1, 5 \cdot 10^5]$, regardless of the integration method. The trapezoidal rule requires almost 10^6 samples to guarantee satisfactory accuracy. Note $\partial_x f$ has a magnitude proportional to K, and thus the variance of the sequence

$$\{\partial_x f(x^1), \partial_x f(x^2), ..., \partial_x f(x^N)\}$$

is of the order of K^2. Therefore, the Monte Carlo approach applied to Eq. 3.9 requires $\mathcal{O}(10^{10})$ samples to secure error of the order of 1. A similar error can be achieved if we perform integration by parts and compute g and generate only $\mathcal{O}(1)$ samples, since the variance is reduced 10^{10} times.

In conclusion, the computational cost of the Monte Carlo method can be dramatically reduced using the generalized fundamental calculus theorem. In case of the f function, the regularization of the integral in Eq. 3.9 may decrease the cost even K^2 times. This result is significant specifically in the context of strongly fluctuating functions. Our approach, based on the integration by parts and systematic computation of g, is particularly designed for problems in which the closed-form expression of the measure itself is unknown. In the example we presented, we numerically reconstruct the measure (and its derivative) by integrating the system in time and applying the derived formulas directly to the generated data. Indeed, data-driven simulations of this type are essential in many computational dynamics algorithms; for example, in sensitivity analysis of chaos [32, 2]. There exist physically inspired systems, however, for which an explicit formula for the measure is known. For example, the statistical behavior of certain systems appearing in the kinetic theory of gases is described by the well-understood Maxwell-Boltzmann speed distribution. In this case, one could directly apply advanced numerical integration techniques dedicated to highly-oscillatory integrands, including Filon- and Levin-based methods [106, 37, 82] and many other; the reader is referred to [134] for a comprehensive review of these methods. Here, we develop a computational tool applicable to generic physics-based and data-driven simulations and, therefore, analysis of different numerical integration rules is beyond the scope of this chapter.

3.2.3 One-dimensional smooth manifolds

We extend the concepts introduced in Section 3.2.1 to the case in which $x(\xi)$ differentiably maps $[0,1]$ to $\mathcal{C} \subset \mathbb{R}^n$, where n is some positive integer. Geometrically, $x(\xi)$ represents a curve embedded in the n-dimensional Euclidean space. The measure $\xi(x)$ can now be expressed as an integral of the density, $\rho : \mathcal{C} \to [0,1]$, along $\mathcal{C}$ with respect to the arc length s,

$$\xi(x) = \int_{\mathcal{C}[x(0),\, x(\xi)]} \rho(x)\; ds, \tag{3.14}$$

where $\mathcal{C}[x(0), x(\xi)]$ denotes a segment of $\mathcal{C}$ between the points indicated in the square bracket. Due to the parameterization $x(\xi)$, the length of the curve $\mathcal{C}$ equals $\int_{\mathcal{C}} ds$, while the arc length differential ds is related to $d\xi$ by $ds = \|dx/d\xi\|\; d\xi$. Using this relation and Eq. 3.14, we obtain the following identity,

$$\rho(x(\xi))\;\; \left\| \frac{dx}{d\xi}(\xi) \right\| = 1. \tag{3.15}$$

We now differentiate Eq. 3.15 with respect to ξ, apply the chain rule and reshuffle terms,

$$g(x(\xi)) = \partial_s \log(\rho(x(\xi))) = \frac{\partial_s \rho}{\rho}(x(\xi)) = -\frac{\frac{dx}{d\xi}(\xi) \cdot \frac{d^2x}{d\xi^2}(\xi)}{\|\frac{dx}{d\xi}(\xi)\|^3}, \tag{3.16}$$

where ∂_s denotes the directional derivative along the curve $\mathcal{C}$ in the direction of increasing ξ. Note the expression for g in Eq. 3.16 reduces to Eq. 3.6 if $x(\xi)$ represents a line manifold, i.e., $\mathcal{C} \subset \mathbb{R}^1$.

As an example, we re-consider the Van der Pol oscillator (Eq. 3.3). This time, however, $x(\xi)$ represents a curve embedded in $\mathbb{R}^2$. In particular, $x(\xi)$ describes a two-dimensional loop such that

$$x(\xi) = \begin{bmatrix} u(\xi T) \\ \frac{du}{dt}(\xi T) \end{bmatrix} \tag{3.17}$$

(see Figure 3-1 for an illustration of the loop). If a numerical solution to Eq. 3.3 is

available, one can combine Eq. 3.15 with Eq. 3.17 to directly evaluate the density function. Similarly, by plugging Eq. 3.17 to Eq. 3.16, it is possible to compute the density gradient, analogously to the procedure described in Section 3.2.1. Consequently, on the RHS of Eq. 3.16, $dx/d\xi$ can be replaced with du/dt, and $d^2x/d\xi^2$ with d^2u/dt^2. We can do so because the density gradient is invariant to any linear transformation of variables. Figure 3-4 illustrates the density function ρ, as well as the length of the curve segment $\mathcal{C}[x(0), x(\xi)]$, versus the parameter ξ. We observe ρ is large if the slope of the length function is small, and vice versa, which is analogous to the $x - \rho$ relation in Figure 3-2. In this particular case, $\rho(\xi)$ is clearly a periodic function with period 0.5. This property is manifested in Figure 3-1. Indeed, one can notice the relation between du/dt and u at $t \in [0, T_{1/2}]$ is the same as $-du/dt$ and $-u$ at $t \in [T_{1/2}, T]$, where $T_{1/2}$ corresponds to $\xi = 0.5$. Figure 3-5 shows the density gradient g computed using two distict ways: through a direct evaluation via Eq. 3.16 and a finite difference scheme (see the caption for more details). The two approaches provide visibly identical solutions, which confirms the correctness of Eq. 3.16. Clearly, the density gradient inherits the periodic behavior of ρ. We notice that the density gradient vanishes if the numerator of Eq. 3.16 is zero, which can happen if $d^2u/dt^2 = 0$ (at the two orange dots in Figure 3-1) and/or $du/dt + d^3u/dt^3 = 0$ (at the six green dots in Figure 3-1). These two cases coincide with the local extrema of the density function (i.e., $d\rho/d\xi = 0$ if at least one of these equations is satisfied). However, zero density gradient does not imply the inflection point ($d^2u/dt^2 = 0$), in contrast to the line manifold case (see Section 3.2.1).

3.3 Computing density gradients on general smooth manifolds

The purpose of this section is to generalize the concept of the density gradient and derive a formula for g defined on higher-dimensional manifolds. Here, we consider a smooth invertible map $x(\xi) : U \to M$, where $U \subset \mathbb{R}^m$, $M \subset \mathbb{R}^n$, $m \leq n$, $x =$

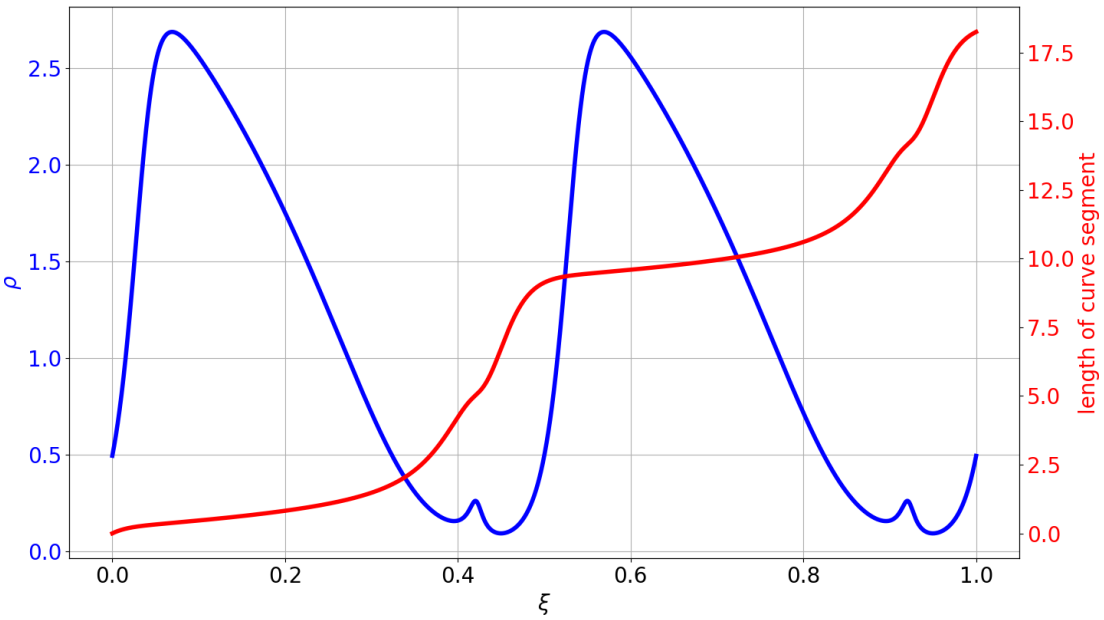

Figure 3-4: The density ρ (blue) and length of the curve segment $\mathcal{C}[x(0), x(\xi)]$ (red) associated with the map $x(\xi)$ defined by Eq. 3.17. The former is computed using the analytical expression in Eq. 3.15, while the latter is approximated by summing the length of consecutive linear segments connecting the points in the sequence $\{x(0), x(\Delta t/T), x(2\Delta t/T), ..., x(\xi)\}$, obtained in the numerical integration of Eq. 3.3.

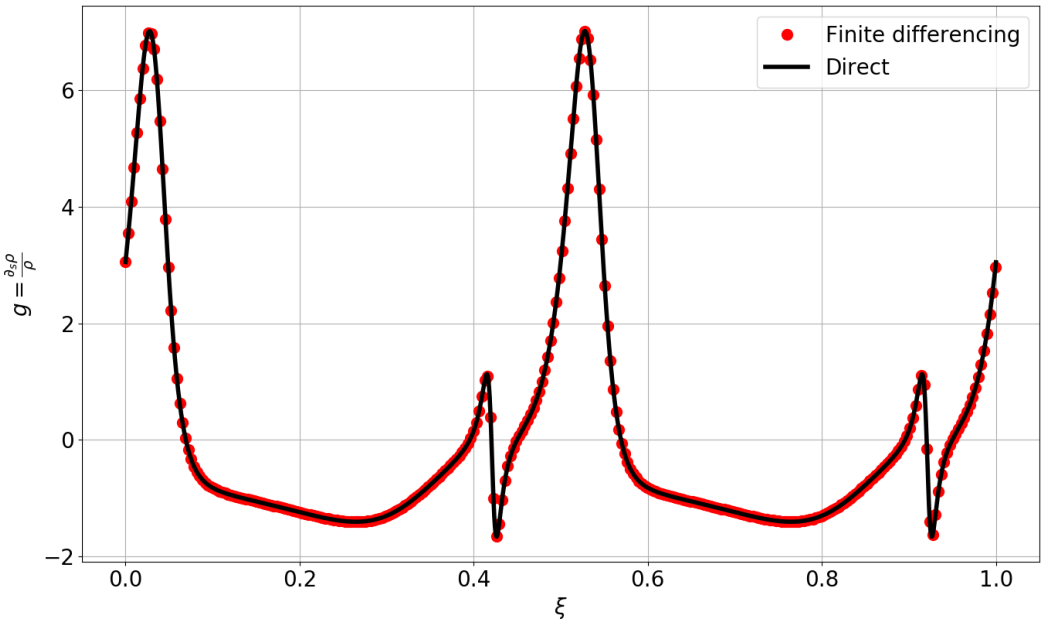

Figure 3-5: The density gradient function g computed directly (using Eq. 3.16) and through a finite difference method. In the latter approach, we note $\partial_s \rho = \partial_\xi \rho / \partial_\xi s$. Both the numerator and denominator is approximated using the central finite difference scheme on a uniform grid using data presented in Figure 3-4.

$[x_1, ..., x_n]^T$ and $\xi = [\xi_1, ..., \xi_m]^T$. U is an m-orthotope (hyperrectangle), which is defined as the Cartesian product of m 1D line manifolds (i.e., intervals of the real line). We no longer assume that these elementary sets only involve numbers between 0 and 1. M is an oriented differentiable manifold, whose shape is defined by the chart x. For example, if $m = 2$ and $n = 3$, then M represents a smooth surface. The density gradient g is now defined as a directional gradient of the logarithm of the density function $\rho : M \to [0, 1]$ implied by the chart $x(\xi)$. In particular, $g = \nabla_s \log \rho$, $\nabla_s := [\partial_{s_1}, \partial_{s_2}, ..., \partial_{s_m}]^T$, where ∂_{s_i}, $i = 1, ..., m$, denote directional derivatives along the corresponding *isoparametric curves*. The i-th component of g is the rate of change of $\log \rho$ along the curve whose preimage involves vectors $\xi \in U$ with constant all coordinates except ξ_i. If $\log \rho$ is differentiable with respect to all the coordinates of x and $\nabla_x := [\partial_{x_1}, \partial_{x_2}, ..., \partial_{x_n}]^T$, then $\partial_{s_i} \log \rho = \nabla_x \log \rho \cdot s_i$, where s_i denotes the unit vector that is tangent to the corresponding isoparametric curve and points in the direction of increasing ξ_i. In Section 3.3.1, we derive a generic formula for g, while Section 3.3.2 provides a specific example of a two-dimensional smooth manifold embedded in $\mathbb{R}^3$ (with $m = 2$ and $n = 3$).

3.3.1 Derivation of the general formula

Recall $x(\xi) : U \to M$ is an invertible and differentiable map, where $U \subset R^m$, $M \subset R^n$, and $m \leq n$, while $\rho(x) : M \to [0, 1]$ is the density function implied by that chart. Let $\omega(x)$ be the natural volume form defined on M. Therefore, the Lebesgue measure m of any subset $V \subset U$, mapped by x to $N \subset M$, equals

$$m(V) = \int_N \rho(x)\, d\omega(x), \tag{3.18}$$

which implies that the measure differential dm defined on U can be expressed in terms of ρ and the volume element defined on M, at every point $x(\xi)$,

$$dm = d\xi_1 \wedge d\xi_2 \wedge ... \wedge d\xi_m = \rho(x)\, d\omega(x) = \rho(x)\, dx_1 \wedge dx_2 \wedge ... \wedge dx_n. \tag{3.19}$$

The wedge symbol ($\wedge$) denotes the exterior product, while $d\xi_i$, $i = 1, ..., m$ and dx_i, $i = 1, ..., n$ represent covectors (1-forms) associated with the corresponding coordinate directions. Intuitively, these 1-forms measure small displacements in the direction of one coordinate. The volume element on M, $d\omega$, can be expressed in terms of ξ,

$$d\omega(x(\xi)) = \sqrt{\det C(x(\xi))}\; d\xi_1 \wedge d\xi_2 \wedge ... \wedge d\xi_m, \tag{3.20}$$

where C represents the $m \times m$ metric tensor of the coordinate transformation $\xi \to x$, defined as

$$C(x(\xi)) = [\nabla_\xi x(\xi)]^T\; \nabla_\xi x(\xi), \tag{3.21}$$

or, componentwise,

$$C_{ij}(x(\xi)) = \partial_{\xi_i} x(\xi) \cdot \partial_{\xi_j} x(\xi). \tag{3.22}$$

The vector gradient $\nabla_\xi x(\xi)$ is represented by an $n \times m$ matrix, in which the j-th column contains the derivative of x with respect to ξ_j, i.e., $[\nabla_\xi x(\xi)]_{ij} = \partial_{\xi_j} x_i(\xi)$. Combining Eq. 3.19 and 3.20, we conclude that the relation between the density function ρ and metric tensor C, at any point $x(\xi) \in M$, can be written in the following way,

$$\rho(x(\xi))\; \sqrt{\det C(x(\xi))} = 1, \tag{3.23}$$

which is a generalization of Eq. 3.15. Let us now QR-factorize the vector gradient $\nabla_\xi x(\xi)$,

$$\nabla_\xi x(\xi) = Q(x(\xi))\; R(x(\xi)), \tag{3.24}$$

where Q is an $n \times m$ matrix, whose columns form an orthonormal basis for the column space of $\nabla_\xi x(\xi)$, while R is an $m \times m$ upper-triangular matrix. Note $Q^T Q = I$ everywhere on M. Using this property, we immediately notice that $C = R^T R$ and, therefore, Eq. 3.23 reduces to

$$\rho(x(\xi))\; |\det R(x(\xi))| = 1. \tag{3.25}$$

For any invertible matrix $A(s)$, which depends on a scalar s, the following indentity is true,

$$\frac{\partial \det A(s)}{\partial s} = \det A \, \mathrm{tr}\left(A^{-1}(s) \, \frac{\partial A(s)}{\partial s}\right). \tag{3.26}$$

Differentiating Eq. 3.25 with respect to ξ_i, applying chain rule and Eq. 3.26, we obtain the following expression for the i-th component of the density gradient,

$$g_i(x(\xi)) = \frac{\partial_{s_i}\rho(x(\xi))}{\rho(x(\xi))} = -\frac{\partial_{s_i} \det R(x(\xi))}{\det R(x(\xi))} = -\frac{\partial_{\xi_i} \det R(x(\xi))}{\det R(x(\xi))\|\partial_{\xi_i} x(\xi)\|}. \tag{3.27}$$

Eq. 3.27 is computationally inconvenient, as it involves evaluating the determinant of R and its directional derivative. Our goal is to rewrite the RHS of that equation such that only first and second parametric derivatives of $x(\xi)$, as well as Q and R factors, are involved.

Since R is an upper-triangular matrix, we notice that

$$\frac{\partial \det R}{\det R} = \frac{\partial \left(\prod_{k=1}^{m} R_{kk}\right)}{\prod_{k=1}^{m} R_{kk}} = \sum_{k=1}^{m} \frac{(\partial R)_{kk}}{R_{kk}} = \mathrm{tr}(\partial R \; R^{-1}). \tag{3.28}$$

Now, differentiating Eq. 3.24 with respect to ξ_i, and then left- and right-multiplying the resulting expression by Q^T and R^{-1}, respectively, we obtain

$$Q^T(x(\xi)) \; \partial_{\xi_i} \nabla_\xi x(\xi) \; R^{-1}(x(\xi)) = Q^T(x(\xi)) \; \partial_{\xi_i} Q(x(\xi)) + \partial_{\xi_i} R(x(\xi)) \; R^{-1}(x(\xi)). \tag{3.29}$$

Note that since $Q^T Q = I$, then $Q^T \; \partial_{\xi_i} Q$ is anti-symmetric, which means its trace vanishes. Therefore, the following equality

$$\mathrm{tr}\left(Q^T(x(\xi)) \; \partial_{\xi_i} \nabla_\xi x(\xi) \; R^{-1}(x(\xi))\right) = \mathrm{tr}\left(\partial_{\xi_i} R(x(\xi)) \; R^{-1}(x(\xi))\right) \tag{3.30}$$

holds everywhere on M. Finally, by combining Eq. 3.27, 3.28 and 3.30, we obtain the general formula for g_i,

$$g_i(x(\xi)) = \partial_{s_i} \log \rho(x(\xi)) = -\frac{\mathrm{tr}\left(Q^T(x(\xi)) \; \partial_{\xi_i} \nabla_\xi x(\xi) \; R^{-1}(x(\xi))\right)}{\|\partial_{\xi_i} x(\xi)\|}, \tag{3.31}$$

which holds everywhere on M for $i = 1, ..., m$. Using Einstein's summation convention, Eq. 3.31 can be rewritten to

$$g_i(x(\xi)) = -\frac{q_j(x(\xi)) \cdot \partial_{\xi_i} \partial_{\xi_k} x(\xi)\ R^{-1}_{kj}(x(\xi))}{\|\partial_{\xi_i} x(\xi)\|}, \tag{3.32}$$

where $q_j(x(\xi))$ denotes the j-th column of $Q(x(\xi))$. Thus, to directly compute the density gradient at any point on a manifold, all first and second derivatives of the chart $x(\xi)$ must be found. In addition, QR factorization of the vector gradient $\nabla_\xi x$ and inversion of the R matrix must be performed. In practice, inverting the triangular matrix R means solving a linear system using the backward substitution method, which requires $\mathcal{O}(m^2)$ operations. Note Eq. 3.32 reduces to Eq. 3.16 if $m = 1$. In the following section, we present an example illustrating some of these quantities. Although Eq. 3.32 is a formula for the derivative in the direction of a isoparametric curve, we can compute derivatives of $\log \rho$ in an arbitrary direction using the distributive law of the dot product.

3.3.2 Example: a surface manifold

As an example of a surface manifold (with $m = 2$ and $n = 3$), let us consider $x(\xi) = u(\xi) = [u_1(\xi), u_2(\xi), u_3(\xi)]^T$, where $\xi = [c, t]^T$, $-5 \leq c \leq 5$, $0 \leq t \leq 0.4$, $u(\xi)|_{t=0} = [c, c, 28]^T$, and $\partial_t u(\xi) = f(u(\xi))$, where f is defined as follows,

$$\begin{aligned}
\partial_t u_1(\xi) &= 10\ (u_2(\xi) - u_1(\xi)), \\
\partial_t u_2(\xi) &= u_1(\xi)\,(28 - u_3(\xi)) - u_2(\xi), \\
\partial_t u_3(\xi) &= u_1(\xi)\,u_2(\xi) - \frac{8}{3} u_3(\xi).
\end{aligned} \tag{3.33}$$

System 3.33 represents the Lorenz '63 oscillator, which is a mathematical model used for atmospheric convection [110]. This system is known to exhibit chaotic behavior. However, we are interested in the solution in a short time interval, such that the trajectories do not intersect and the resulting surface is orientable. In particular, we compute $x(\xi)$ by numerically integrating System 3.33 in time for different values of

$c \in [-5, 5]$, using the second-order Runge-Kutta scheme with $\Delta t = 0.002$. There are two reasons we have chosen this particular $x(\xi)$. First, it serves as a perfect example of a problem, in which the smooth one-to-one solution, $x(\xi)$, cannot be found analytically. Thus, the computation of g should be performed numerically using closed-form relations derived in Section 3.3.1. Second, the surface described by the chart $x(\xi)$ can be obtained as a evolution of 1D manifolds. This observation is utilized in Section 3.4, where we derive expressions for evolving manifolds. To evaluate ρ and g, we directly use Eq. 3.25 and Eq. 3.31, respectively. To find these quantities, the vector gradient $\nabla_\xi x(\xi) = [\partial_c x(\xi), \partial_t x(\xi)]$, as well as the following second derivatives: $\partial_c^2 x(\xi), \partial_t^2 x(\xi), \partial_c \partial_t x(\xi)$, must be found at every point on the manifold. The time derivative, $\partial_t x(\xi) = f(x(\xi))$, is obtained automatically as we integrate System 3.33 in time. The second derivative of x with respect to t is obtained using the chain rule,

$$\partial_t^2 x(\xi) = \partial_t f(x(\xi)) = Df(x(\xi))\, f(x(\xi)),$$

where Df denotes the Jacobian of System 3.33. Thus, from the computational point of view, we need to solve a tangent equation to find $\partial_t x(\xi)$ at every point of the trajectory defined by System 3.33. Using this approach, one can analogously find derivatives with respect to c. Let $v(\xi) = \partial_c x(\xi)$ and $w(\xi) = \partial_c^2 x(\xi)$. Using the chain rule, we conclude that

$$\partial_t v(\xi) = Df(x(\xi))\, v(\xi),$$

$v(\xi)|_{t=0} = [1, 1, 0]^T$ and, by differentiating again,

$$\partial_t w(\xi) = D^2 f(x(\xi))(w(\xi), w(\xi)) + Df(x(\xi))\, w(\xi),$$

$w(\xi)|_{t=0} = [0, 0, 0]^T$, where $D^2 f$ denotes the Hessian of f. Using Einstein's summation convention, the i-th component of the bilinear form $D^2 f(x(\xi))(w(\xi), w(\xi))$ can be written as $\partial_{x_k} \partial_{x_l} f_i\, w_k\, w_l$. Finally, the mixed derivative $\partial_c \partial_t x(\xi) = \partial_t v(\xi)$ is a byproduct of the numerical integration of the tangent equation for v. We solve all of these tangent equations using the same time integrator as the one mentioned

above. Since $m = 2$, the 2×2 R matrix is inverted analytically at every point on the trajectory.

In this case, the U space, which is the domain (preimage) of x, is in fact a Cartesian product of $[-5, 5]$ and $[0, 0.4]$. The upper plot in Figure 3-6 graphically represents U, while the lower plot illustrates the $u_1 - u_3$ projection of M, obtained through the mapping $x(\xi)$. For completeness, in Figure 3-7, we also include the $u_1 - u_2$ and $u_2 - u_3$ projection of the deformed mesh. It is clear that the deformation is symmetric with respect to $c = 0$. We also observe that fibers (isoparametric lines) corresponding to larger values of t are subject to greater stretching than those at smaller t. These features are reflected by the distribution of the density function ρ, plotted in Figure 3-9. The smaller the area of each distorted quadrilateral of the mesh, the larger the value of the density function. Indeed, the smallest values of the density distribution are located around $t = 0.4$. This region coincides with the most stretched quadrilaterals.

Figure 3-8 shows the two components of the density gradient $g := [g_c, g_t]^T = [\partial_{s_1} \log \rho, \partial_{s_2} \log \rho]^T$, corresponding respectively to the c- and t-direction. The distribution of g_c is clearly symmetric with respect to the reflection points on the isoparametric line $c = 0$, which is a manifestation of the fact the density is symmetric and directional derivative is computed in the direction of increasing c. The symmetry of g_t is a direct consequence of the definition $g_t := \partial_{s_2} \log \rho$, where $\log \rho$ itself is symmetric. Note the largest-in-magnitude values of g_a concentrate around the boundaries of the range of c, i.e., at $c = \pm 5$ and, in case of g_t, around $u_1 = 0$. This reflects the fact the density gradient measures the relative rate of change of the density. In particular, its value becomes large if the rate of change of the density is large and/or the density itself is small. Figure 3-10 illustrates the density gradient along the bold isoparametric curves from Figure 3-6, computed using Eq. 3.32 directly and through finite differencing. In case of both g_c and g_t, we observe a good agreement between the solution computed directly and finite difference approximation, which validates our derivation of Eq. 3.32.

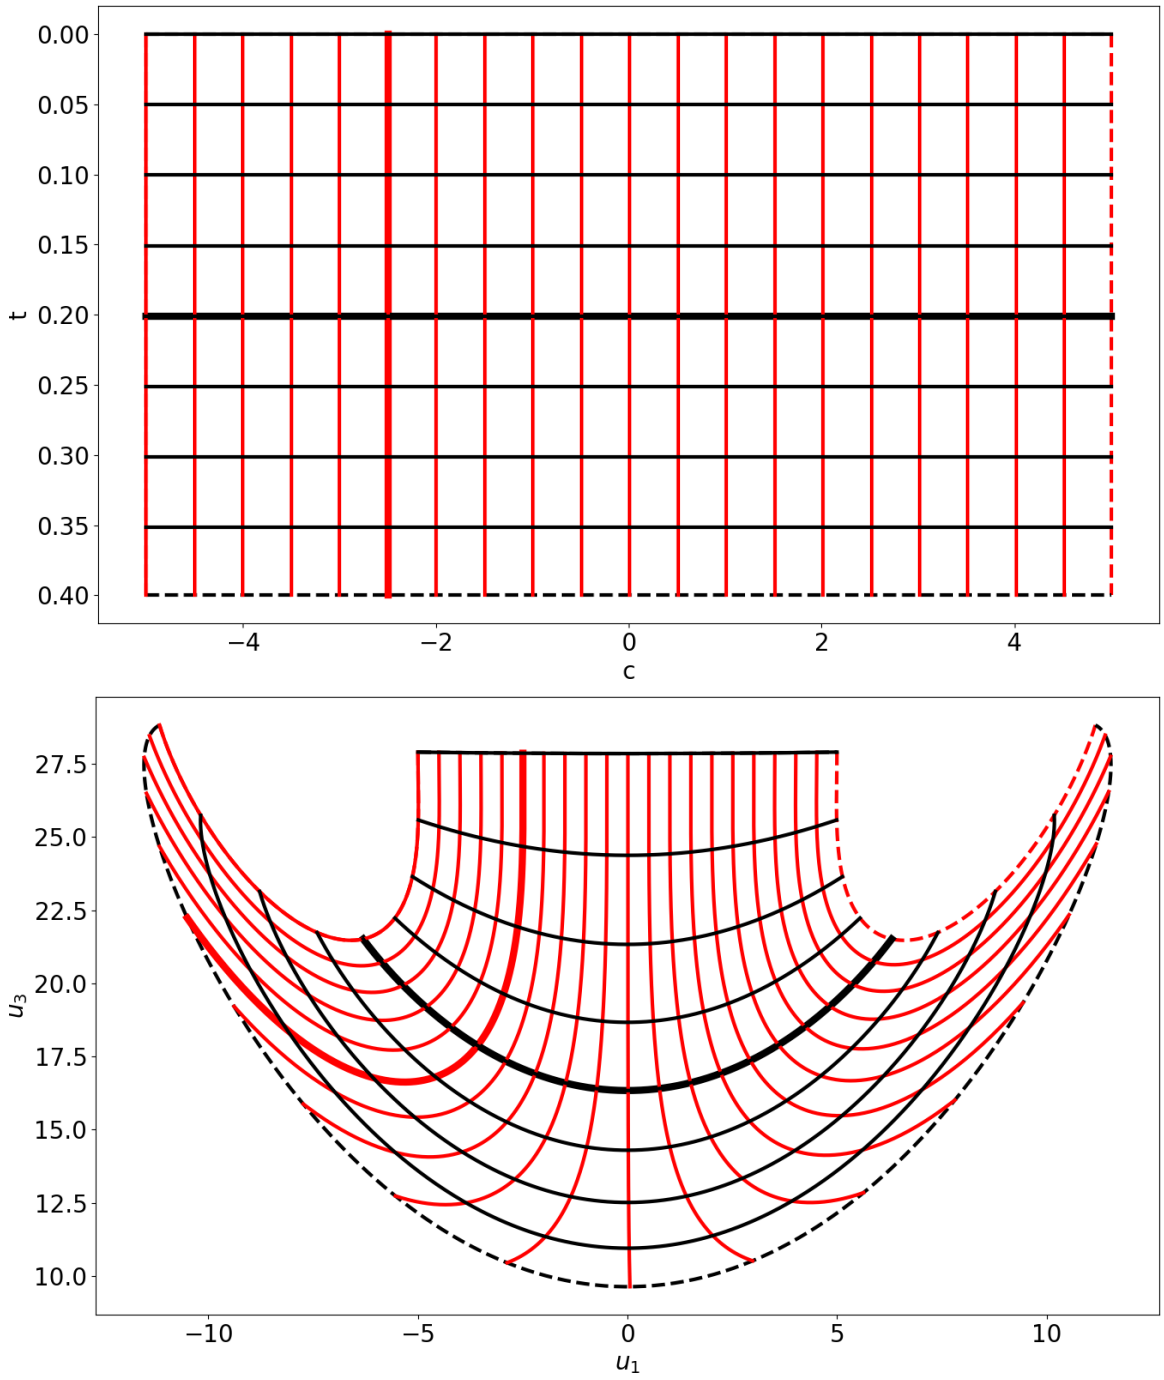

Figure 3-6: Upper plot: a structured mesh representing the domain $U = \{(c,t) \,|\, c \in [-5,5],\, t \in [0,0.4]\}$. The black lines correspond to fixed values of t, while the red lines illustrate ξ with a fixed value of c. The red and black dashed lines represent $c = 5$ and $t = 0.4$, while the red and black bold lines refer to $c = -2.5$ and $t = 0.2$, respectively. Lower plot: $u_1 - u_3$ projection of the image of the structured mesh obtained through the mapping $x(\xi)$.

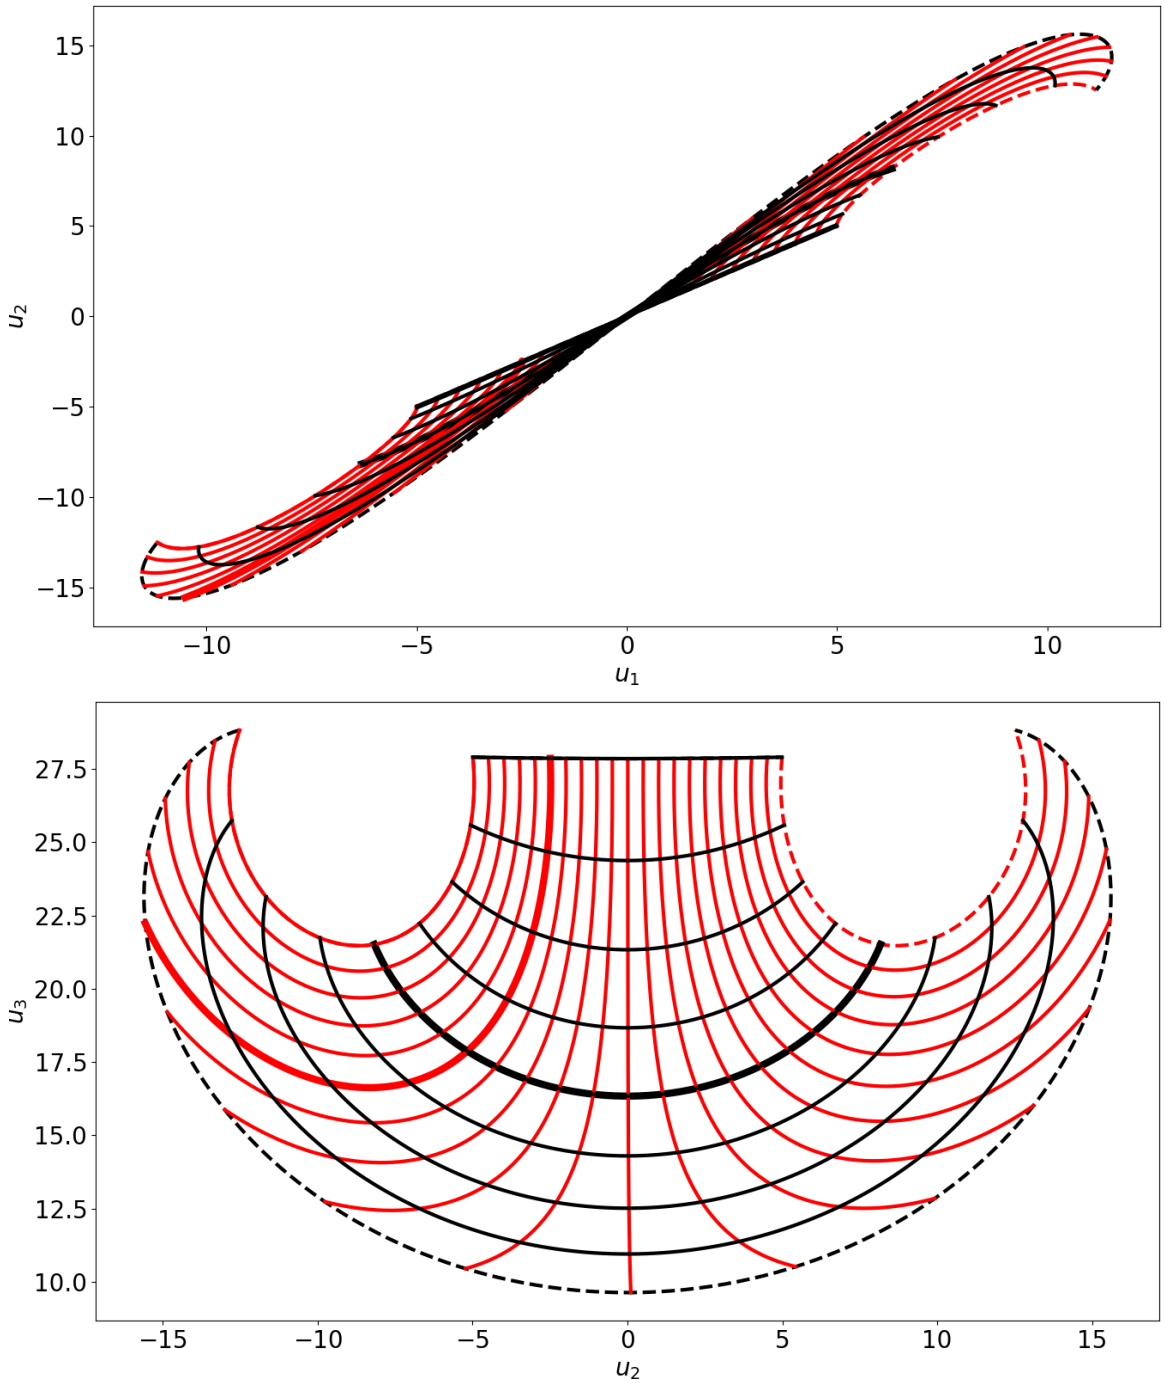

Figure 3-7: Extension of Figure 3-6. $u_1 - u_2$ (upper plot) and $u_2 - u_3$ (lower plot) projection of the image of the structured mesh obtained through the mapping $x(\xi)$.

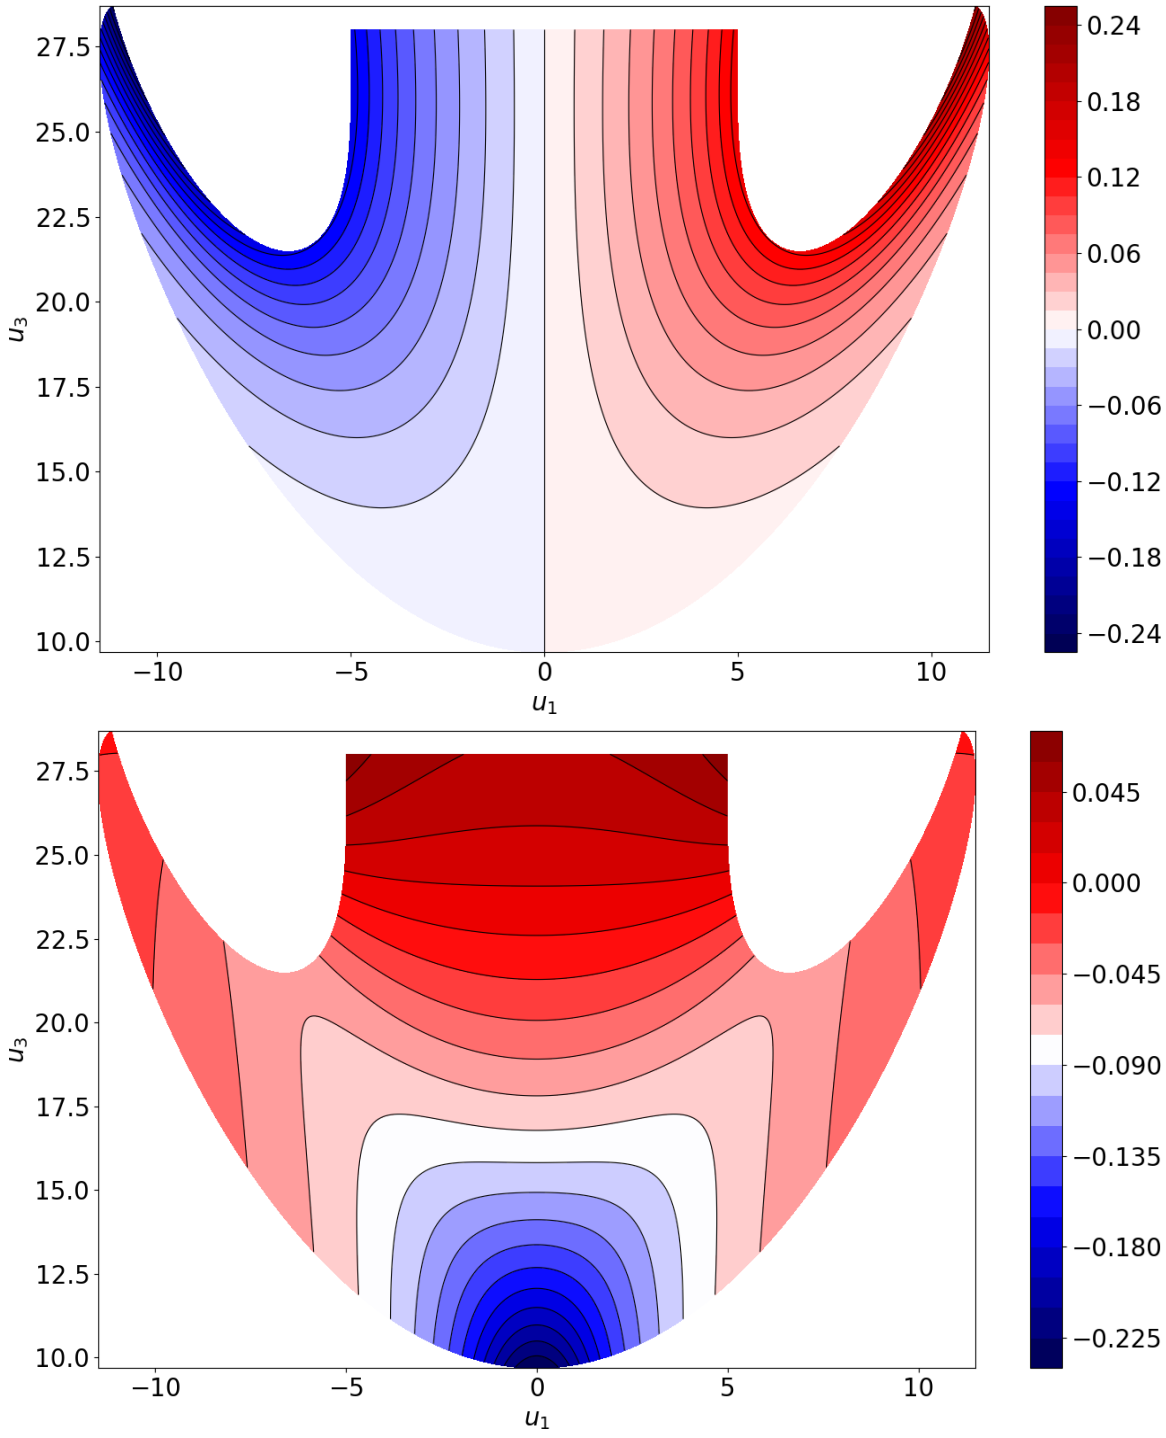

Figure 3-8: $u_1 - u_3$ projection of the directional derivative of $\log \rho$, in the c-direction, $g_1 := g_c$ (upper plot), and t-direction, $g_2 := g_t$ (lower plot).

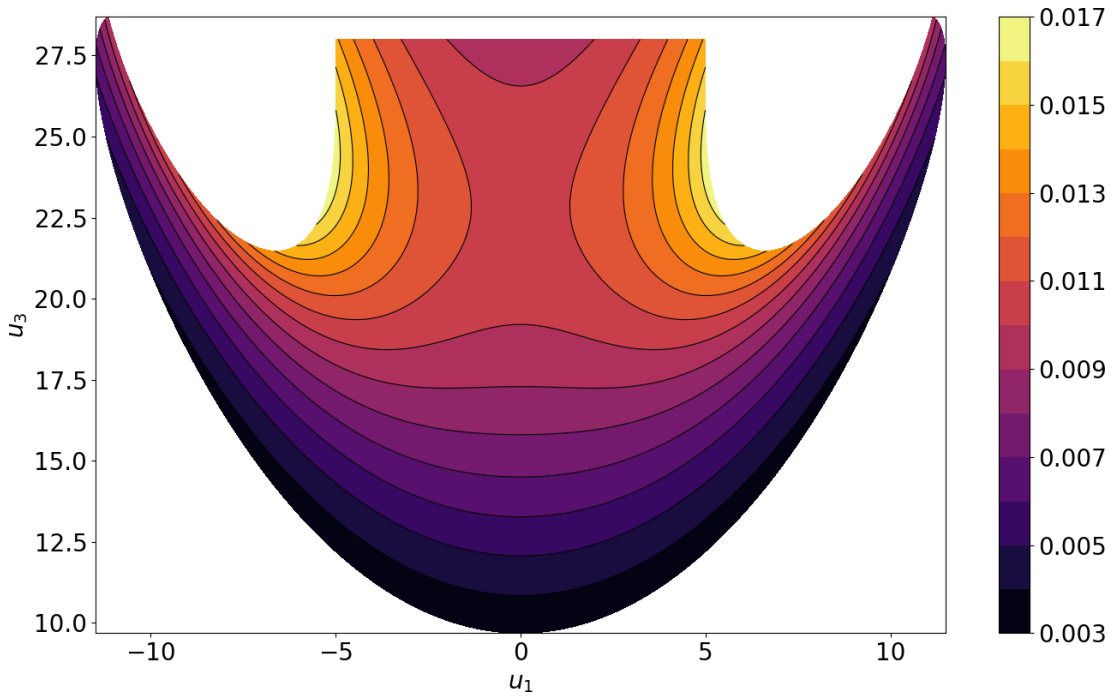

Figure 3-9: $u_1 - u_3$ projection of the density distribution.

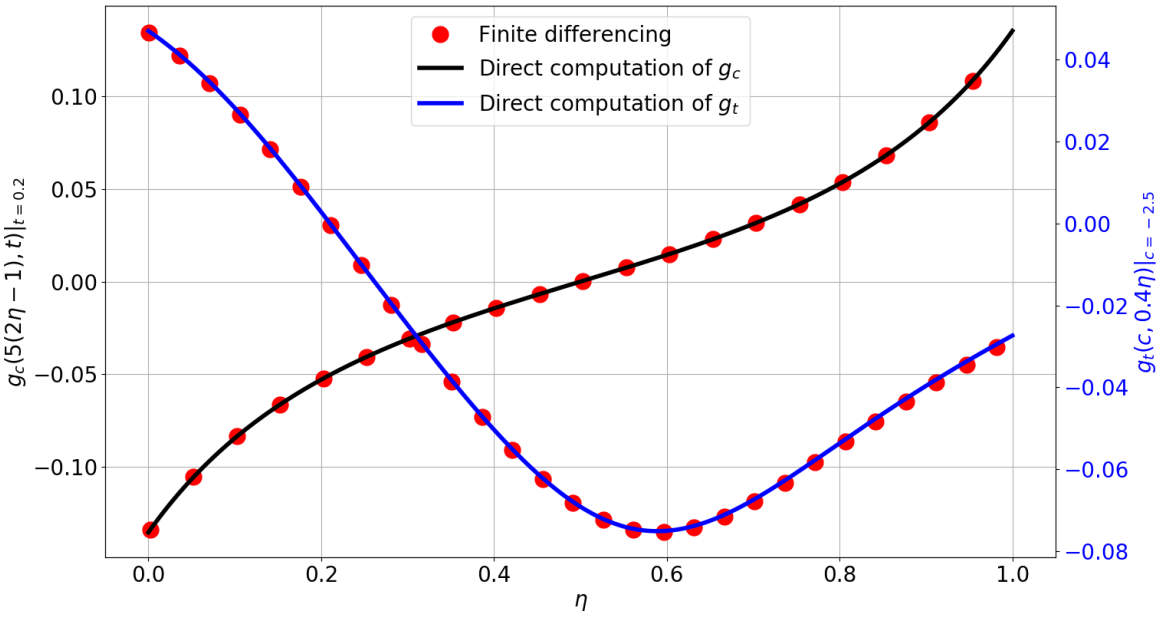

Figure 3-10: The first and second component of the density gradient function $g = [g_c, g_t]$, respectively at $t = 0.2$ (black solid line on Figure 3-6) and $c = -2.5$ (red solid line on Figure 3-6), computed directly using Eq. 3.32 and through a finite difference method. In the latter approach, we note $\partial_{s_i}\rho = \partial_{\xi_i}\rho/\partial_{\xi_i}s_i$, $i = 1, 2$, $\xi_1 = c$, $\xi_2 = t$, where s_i denotes the length of the isoparametric curve associated with ξ_i. Both the numerator and denominator is approximated using the central finite difference scheme on a uniform grid. The relation $s_i(\xi_i)$ is found in a way analogous to the one described in Section 3.2.3.

3.4 Recursive computation for density gradients along trajectories implied by multi-dimensional non-linear transformations

Using the results presented in Section 3.2 and 3.3, we now propose an iterative method for the density gradient along trajectories defined by a C^2 diffeomorphism $\varphi : M^k \to M^{k+1}$, $k \in \mathbb{Z}$, where both M^k and M^{k+1} represent differentiable manifolds of the same dimension embedded in $\mathbb{R}^n$, $n \in \mathbb{Z}^+$. Let us consider two different charts, $x^k(\xi) \in N^k \subset M^k$ and $x^{k+1}(\xi) \in N^{k+1} \subset M^{k+1}$, such that

$$x^{k+1}(\xi) = \varphi(x^k(\xi)) \tag{3.34}$$

for all $\xi \in V \subset U \subset \mathbb{R}^m$, $1 \leq m \leq n$, $k \in \mathbb{Z}$. Let ω^k and ω^{k+1} be the natural volume forms in M^k and M^{k+1}, respectively, where ω^{k+1} is the pushforward of ω^k under φ. Therefore, for all $k \in \mathbb{Z}$, the Lebesgue measure m of the subspace V can be expressed as follows,

$$m(V) = \int_{N^k} \rho^k(x)\, d\omega^k(x) = \int_{N^{k+1}} \rho^{k+1}(x)\, d\omega^{k+1}(x), \tag{3.35}$$

where ρ^k and ρ^{k+1} are densities implied by $x^k(\xi)$ and $x^{k+1}(\xi)$, respectively. Following the procedure involving Eq. 3.19-3.21, it is possible to find the relation between ρ^k, ρ^{k+1}, and the metric tensors of the two transformations: $\xi \to x^k$ and $\xi \to x^{k+1}$. Thus, by applying the chain rule, we find a relation between the parametric derivatives of $x^k(\xi)$ and $x^{k+1}(\xi)$, thanks to which a general recursive formula for the density gradient along the trajectory defined by φ can be inferred. The g^k function should be understood as the directional derivative of the (logarithmic) density implied by the chart $x^k(\xi)$. In Section 3.4.1, we derive an iterative procedure for g^k, while Section 3.4.2 presents the use of the proposed algorithm by revisiting the Lorenz '63 oscillator. Throughout this section, repeated indices in the subscript of any term imply summation (Einstein's notation), unless otherwise stated.

3.4.1 A generic recursive procedure for evolving density gradients

As pointed out above, the first step is to find a relation between the parametric gradients of x^k and x^{k+1}. Applying the definition of φ from Eq. 3.34 and the chain rule, we can expand $\nabla_\xi x^{k+1}$ in the following way,

$$\nabla_\xi x^{k+1}(\xi) = D\varphi(x^k(\xi))\ \nabla_\xi x^k(\xi), \tag{3.36}$$

or, equivalently,

$$\partial_{\xi_i} x^{k+1}(\xi) = D\varphi(x^k(\xi))\ \partial_{\xi_i} x^k(\xi), \tag{3.37}$$

where $D\varphi$ denotes the $n \times n$ Jacobian matrix of φ, i.e., $(D\varphi)_{ij} = \partial_{x_j}\varphi_i$. By differentiating Eq. 3.37 once more, with respect to ξ_j, we obtain

$$\partial_{\xi_i}\partial_{\xi_j} x^{k+1}(\xi) = D^2\varphi(x^k(\xi))\left(\partial_{\xi_i} x^k(\xi), \partial_{\xi_j} x^k(\xi)\right) + D\varphi(x^k(\xi))\ \partial_{\xi_i}\partial_{\xi_j} x^k(\xi), \tag{3.38}$$

where $D^2\varphi$ is the Hessian of φ, which is in fact a third-order $n \times n \times n$ tensor. Analogously to the example presented in Section 3.3.2, the first term in the RHS of Eq. 3.38 is a bilinear form that outputs an n-element vector. In this case, the i-th component of that vector equals

$$\partial_{x_p}\partial_{x_q}\varphi_i(x^k(\xi))\ \partial_{\xi_i} x_p^k(\xi)\ \partial_{\xi_j} x_q^k(\xi).$$

In the second step, we directly use the formula for the density gradient derived in Section 3.3.1. Let $f^k := f(x^k(\xi))$ be a shorthand notation for any function f defined at $x^k(\xi)$, and $e_i(x^k(\xi)) := \partial_{\xi_i} x^k(\xi)$, $a_{ij}(x^k(\xi)) := \partial_{\xi_i}\partial_{\xi_j} x^k(\xi)$. Thus, by combining Eq. 3.37, 3.38 with Eq. 3.32 derived for a generic chart $x(\xi)$, we conclude that

$$g_i^k = -\frac{(R_{lj}^{-1})^k}{\|e_i^k\|}\, q_j^k \cdot a_{il}^k, \tag{3.39}$$

$$(\nabla_\xi x)^k = [e_1^k\, e_2^k \cdots e_m^k] = Q^k\, R^k = [q_1^k\, q_2^k \cdots q_m^k]\, R^k. \tag{3.40}$$

$$e_i^{k+1} = D\varphi^k e_i^k, \tag{3.41}$$

$$a_{ij}^{k+1} = D^2\varphi^k(e_i^k, e_j^k) + D\varphi^k a_{ij}^k, \tag{3.42}$$

hold for any $\xi \in V \subset U$.

To summarize, if a map φ relating two consecutive points on the trajectory, $x_k(\xi)$ and $x_{k+1}(\xi)$, is available, then the density gradient at one point can be computed using information associated with the other point. In particular, according to Eq. 3.39, the i-th component of g requires knowledge of e_j, $j = 1, ..., m$ and a_{pq}, $p, q = 1, ..., m$ at the same point. Thus, to compute one component of the density gradient at $x^k(\xi)$ for some ξ, we need to apply the recursion in Eq. 3.41 km times and, analogously, the recursion in Eq. 3.42 $1/2\,km^2$ times. The $1/2$ factor is a consequence of the fact that $a_{ij}(\xi) = a_{ji}(\xi)$ for any admissible ξ, because x_k is assumed to be twice differentiable for any $k \in \mathbb{Z}$. In addition, at every step k, the QR factorization of $(\nabla_\xi x)^k = [e_1^k\, e_2^k \cdots e_m^k]$ and inversion (either direct if m is small or through solving a linear system) of the resulting $m \times m$ R^k matrix must be performed. We assume $x^0(\xi)$ is given, from which we directly compute initial conditions for recursions in Eq. 3.41 and 3.42.

The recursion involving Eq. 3.39-3.42 can be used to devise algorithms for differentiating the invariant, physical SRB measure m_{SRB}, which is guaranteed to exist in uniformly hyperbolic systems. The SRB measure is in fact the pushforward of the Lebesgue measure on unstable manifolds; it contains a statistical description of the dynamics. In general, m_{SRB} is not absolutely continuous everywhere on the manifold, but only conditional measures of m_{SRB} along unstable manifolds are absolutely continuous. The SRB density gradient g_{SRB}, defined as a directional derivative of the conditional SRB density on the unstable manifold, is a byproduct of the integration by parts (analogous to Eq. 3.2), preceded by the disintegration of m_{SRB} [32, 158]. Thus, if a direction of the unstable manifold is given, the recursive formula presented in this section might be further developed to compute the SRB density gradient, defined on a manifold of any dimension, along a trajectory initiated at a m_{SRB}-typical point.

3.4.2 Example: evolution of a 1D manifold

In this section, we demonstrate the application of the recursive scheme for the density gradient g^k. For this purpose, let us re-consider the Lorenz '63 oscillator, defined by System 3.33. In particular, we define φ, such that it represents numerical time integration of System 3.33 for a period of Δt, i.e., $u(t+\Delta t) = \varphi(u(t))$ with $u(t)$ being the solution of the system at time t. Let us consider a 1D smooth manifold embedded in $\mathbb{R}^3$ described by the following chart $x^0(c) = [c, c, 28]^T$, $-5 \leq c \leq 5$. Note $x^0(c)$ coincides with the black solid boundary of the surface depicted in Figures 3-6-3-7. Now, by applying φ recursively, the next step is to numerically compute a sequence of charts $\{x^0(c), x^1(c), x^2(c), ...\}$, where $x^{k+1}(c) = \varphi(x^k(c))$. Our aim is to compute $g^k = \partial_s \log \rho^k$, where ρ^k is a density implied by the chart $x^k(c)$. The operator ∂_s denotes a generic directional derivative along the curve in the direction of increasing c. The formulas derived in Section 3.4.1 give us all necessary tools to compute g^k along the trajectory defined by φ. In this example, however, we consider the simplest case, $m = 1$. Eq. 3.39-3.42 can be dramatically simplified, because $\nabla_\xi x = dx/dc$ is just a vector, and thus QR factorization is equivalent to normalizing that vector. Let $e = dx/dc = \|dx/dc\|\, q$ and $a = d^2x/dc^2$ and, therefore,

$$g^k = -\frac{q^k \cdot a^k}{\|e^k\|^2}, \tag{3.43}$$

$$e^{k+1} = D\varphi^k\, e^k, \quad q^k = \frac{e^k}{\|e^k\|}, \tag{3.44}$$

$$a^{k+1} = D^2\varphi^k(e^k, e^k) + D\varphi^k\, a^k. \tag{3.45}$$

Note Eq. 3.43-3.45 can be derived directly using Eq. 3.16 and the chain rule for parametric derivatives.

How does this example differ from the one presented in Section 3.3.2? There, we used a chart $x_s(\xi) : \mathbb{R}^2 \to \mathbb{R}^3$, $\xi = [c, t]^T$, which defined a two-dimensional manifold. The rate of change of $x(\xi)$ in the t-direction was determined by the Lorenz '63 oscillator (System 3.33). Here, using the iterative procedure, we generate a bunch of 1D manifolds $x^k(c)$. The evolution of these curves (in geometric sense) is deter-

mined by φ, which is in fact a discrete version of System 3.33. Thus, if we generate infinitely many such curves and $\Delta t \to 0$, we effectively obtain the same surface as the one shown in Figure 3-6. Intuitively, the density ρ_s implied by $x_s(\xi)$ measures number of points mapped from a uniform distribution per unit surface area. Likewise, the density ρ^k, implied by the chart $x^k(c)$, measures number of points mapped from a uniform distribution per unit curve length. Since we use the same discretization scheme to integrate differential equations, the localization of points obtained in both the computation of surface from Section 3.3.2 and, here, evolution of curves is exactly the same. However, the density ρ^k does not equal the marginal distribution of ρ_s at $t = k\,\Delta t$ (assuming uniform discretization of time). In case of the surface example, the value of the density function reflects the densification of points, mapped from a uniform distribution, in both the t and c directions. In the latter example, the density is determined only by the localization of points along the evolving curve. Figure 3-11 illustrates the density gradient g^k along the evolving curve, recorded at three different time steps k. We observe $g^k = 0$ at $k = 0$, which is a consequence of the choice of the uniformly distributed initial condition. Due to the symmetric geometry of M^k, defined by the Lorenz '63 oscillator at $t \in [0, 0.4]$, the density gradient features symmetric behavior with respect to the origin of the $g^k(c)$-vs.-c relation.

3.5 Density gradients implied by coordinate charts: conclusion

A Monte Carlo integration scheme applied to a highly-oscillatory function might be remarkably expensive. The computational cost, however, can be dramatically reduced by integrating the original formulation by parts. Such treatment gives rise to a new quantity, i.e., a directional derivative of the logarithm of the density implied by a chart describing the integration domain. The computation of that derivative, which we call the *density gradient*, requires knowledge of the first and second derivatives of the chart with respect to the domain parameterization. If the domain manifold evolves

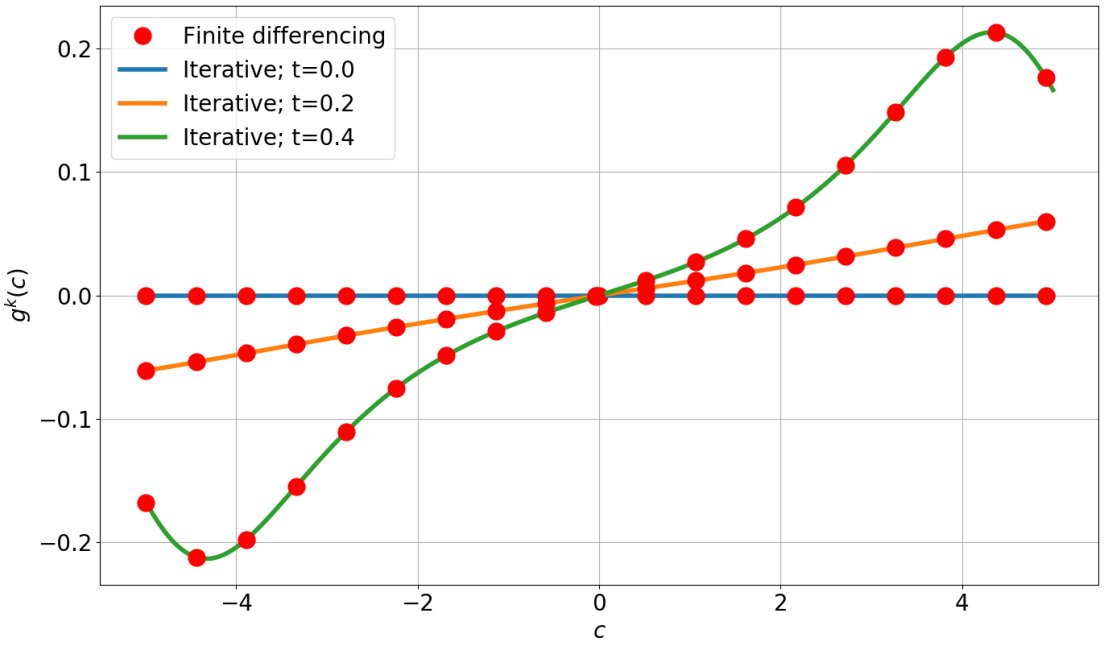

Figure 3-11: Density gradient function is computed using the recursion involving Eq. 3.43-3.45 at three different time steps $k = t/\Delta t$. The finite difference approximation is generated using the approach described in Section 3.2.3.

according to some diffeomorphism φ, the calculation of the density gradient along a trajectory requires solving a collection of first- and second-order tangent equations, involving both the Jacobian and Hessian of φ. The number of these equations is respectively proportional to m and m^2, where m is the dimension of the manifold. The formulas derived above is a major step toward constructing generalizable algorithms for *SRB density gradients*.

3.6 SRB measure and its gradient: significance and definitions

Due to their seemingly irregular and quasi-random behavior, a mathematical description of chaotic dynamical systems might be challenging. A major breakthrough in the analysis of chaos was the introduction of the SRB (Sinai-Ruelle-Bowen) measure μ [147]. This scalar quantity, defined on a compact Riemannian manifold, contains a coherent statistical description of the dynamics. Intuitively, the long-time average

of a smooth observable, computed along a trajectory initiated at any point from a non-zero Lebesgue measure set, converges to the μ-weighted mean of the same observable. Although the concept of SRB measures was originally applied to Axiom A systems, several rigorous studies extended this idea beyond the universe of uniformly hyperbolic systems [195, 40, 42, 153].

Lebesgue integrals with respect to μ, which indeed represent expected values of certain smooth observables, are fundamental in the analysis of chaos. Under the assumption of *ergodicity*, they equal the time-average of an infinitely-long sequence generated along a trajectory. Integrals of this type can thus be approximated using a Monte Carlo method. If the integrand involves highly-oscillatory derivatives, then the Monte Carlo integration might be prohibitively expensive due to a large variance of the sample [160]. In case of derivatives of functions evaluated at a future time (see examples of such integrands in [31, 158, 68, 2]), the direct use of any integration scheme might be impossible due to the *butterfly effect*. Indeed, the application of the chain rule results in a product of the system's Jacobian matrices whose norms increase exponentially in time. A remedy for this computational difficulty is integration by parts, which moves the differentiation operator away from the problematic function to the SRB measure. This is in fact a consequence of the generalized fundamental theorem of calculus. In addition to the boundary term, we effectively obtain a new Lebesgue integral involving a product of the antiderivative of the original integrand and the *SRB density gradient* $g = \partial \log \rho = \partial \rho / \rho$, where ρ denotes the density of μ (i.e., the Radon-Nikodym derivative [123]).

The SRB density gradient is critical in the sensitivity analysis of chaos. The major implication of Ruelle's *linear response* theory is a closed-form expression for the parametric derivative of long-time averages (a.k.a. the system's sensitivity) [148, 150]. The space-split sensitivity (S3) method [32, 31, 158] reformulates Ruelle's formula to a computable form by splitting the perturbation vector and performing integration by parts on unstable manifolds. Using the S3 formula, one can construct an efficient and provably convergent Monte Carlo algorithm for sensitivities in uniformly hyperbolic systems. This algorithm requires computing the SRB density gradient

defined as a directional derivative of ρ conditioned on the unstable manifold. Indeed, the SRB measure is generally singular with respect to Lebesgue measure in the stable direction [195]. Several algorithms for sensitivity analysis that stem from the Fluctuation-Dissipation Theorem (FDT) [88] also require g [2, 1, 23, 104, 68]. An accurate reconstruction of the linear response operator in FDT-based methods is the major challenge in deriving reliable numerical schemes. Motivated by empirical data of certain chaotic models, some methods of this type assume Gaussian distribution of measure (see [68] and references therein). Such an assumption reduces the FDT linear response operator to a simple time autocorrelation function, which dramatically decreases the total cost. However, this simplification restricts the algorithm to a narrow class of atmospheric chaotic systems. The density gradient can also be used as an reliable indicator of the differentiability of statistical quantities [159] in chaotic systems. In particular, the slope of the distribution tail of g has been shown to be strictly associated with the existence of parametric derivatives of statistics. Therefore, we seek a direct numerical procedure for g that does not make any assumptions about the statistical behavior of the system and is thus applicable to any chaotic dynamical system that admits SRB measures.

There already exist algorithms for the SRB density gradient derived for systems with one-dimensional unstable manifolds. In case of simple one-dimensional maps, one can derive an exponentially convergent recursion for g using the measure preservation property [158] (see Section 2.6.3). The same formula can be inferred using the fact the SRB density is an eigenfunction of the Frobenius-Perron operator with eigenvalue 1 [159]. The authors of [33] propose an ergodic-averaging algorithm for self-derivatives (i.e., directional derivatives along one-dimensional expanding directions) of covariant Lyapunov vectors (CLVs) corresponding to the only positive Lyapunov exponent, which are tangent to unstable manifolds at any point on the attractor. Using the chain rule on smooth manifolds, one can show g depends on the self-derivative of CLV at the previous time step, and this relation is governed by a second-order tangent equation [33, 159]. In a recent work on a new method for evaluating the linear response [127], the direct computation of g is circumvented by re-expressing the unstable divergence

as the volume ratio between the projection onto the contracting (stable) subspace and the imposed perturbation. Nevertheless, several numerical methods based on the linear response theory still require g, which justifies the need for a generalizable algorithm for this particular quantity.

In the remainder of this chapter, we systematically derive a trajectory-driven algorithm for the SRB density gradient by extending the measure preservation property to high-dimensional smooth manifolds. Using the density-based parameterization of unstable manifolds and the chain rule, it is possible to establish a recursive relation for the evolution of first- and second-order parametric derivatives of the coordinate chart. By definition, this chart is strictly associated with g and can be interpreted as an SRB inverse cumulative distribution (quantile function). This type of parameterization, motivated by popular methods of statistical inference [57], has been thoroughly explained in the first part of this chapter (Sections 3.1 – 3.5) in the context of simple Lebesgue measures. Through the relation of g, the coordinate map and its parametric derivatives, we show the density gradient can be computed by solving a collection of first- and second-order tangent equations. We also show that the recurring problem of the butterfly effect, which leads to exploding norms of tangent solutions, can be eliminated by iterative orthonormalization of the chart gradient. The major benefit of our derivation is that it is naturally translatable to a practicable algorithm that can be easily integrated with existing methods for sensitivity analysis of chaos.

Consider a diffeomorphic map $\varphi : M \to M$, $M \in \mathbb{R}^n$, $n \in \mathbb{Z}^+$ with an Axiom A attractor. Theorem 1 of [195] asserts that there exists an *invariant* and *physical* probability measure μ (and its density ρ), which satisfies:

1. Invariance/conservation of measure condition:

$$\mu(A) = \mu(\varphi^{-1}(A)) \tag{3.46}$$

   for any Borel subset $A \subset M$.

2. Physicality condition: there exists a positive Lebesgue measure set V such that

for any smooth observable $f : M \to \mathbb{R}$,

$$\int_M f(x)\, d\mu(x) = \int_M f(x)\, \rho(x)\, d\omega(x) = \lim_{N\to\infty} \frac{1}{N} \sum_{k=0}^{N-1} f \circ \varphi^k(x_0), \tag{3.47}$$

for all $x_0 \subset V$. We use $d\omega$ to denote the Riemmanian volume element, and $\varphi^k(\cdot) = \varphi(\varphi^{k-1}(\cdot))$, $\varphi^1 = \varphi$, $\varphi^0 = \mathrm{Id}$.

3. Absolute continuity: Conditional measure of μ denoted by $\tilde{\mu}_x$ and defined on the unstable manifold U_x at point $x \in U_x$ is absolutely continuous (an analogous property applies to the conditional density $\tilde{\rho}_x$).

4. Singularity with respect to Lebesgue measure: μ is generally sharp in the stable direction (across unstable manifolds).

5. Unit measure axiom (probability universe):

$$\int_M d\mu(x) = \int_M \rho(x)\ d\omega(x) = \mu(M) = 1.$$

The measure μ and its density ρ are respectively known as the SRB measure and SRB density distribution. We listed their properties most important in the context of this chapter; however, the reader is referred to [195] for a detailed description of other significant features. One can think about Property (1) as the mass conservation law. For example, consider a Borel subset $B \subset M$ with a uniform measure that is mapped to $\varphi(B) \subset M$. If we divide $\varphi(B)$ into a finite number of subsets occupying the same volume, each of them generally has a different measure. In other words, each subset generally has its unique weight unless φ represents a simple translation and/or rotation. Property (2) states that the SRB measure is physical, which means that the system can be observed due to the positive Lebesgue measure sets. Consequently, by "observing" the system's evolution for an infinitely long period of time, we can assign a weight (density) to each non-zero-volume region of the attractor. The expected value of any smooth function defined on M can be computed as a simple volume integral over M of that function weighted by the density function. Figure 3-12 graphically

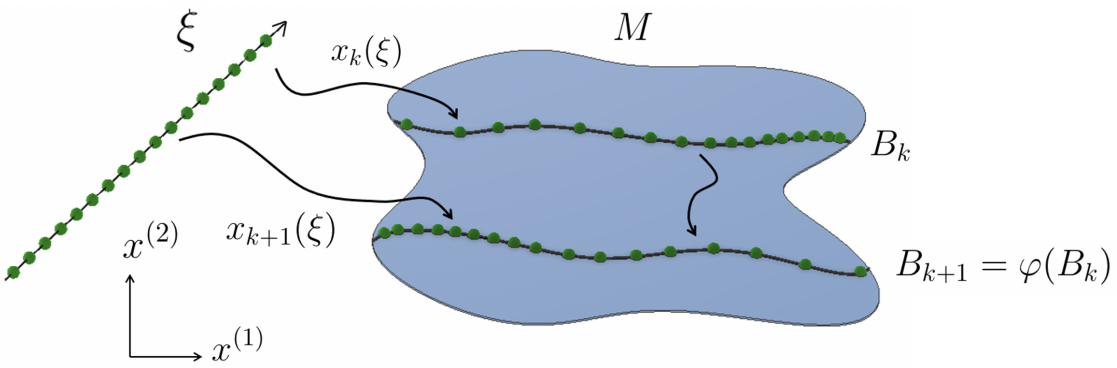

Figure 3-12: This figure graphically represents the measure preservation property. The localization of green bullets represents the SRB density on some 1D subspace of a 2D manifold M. All green bullets are equally weighted. In this sketch, we observe $\mu(B_k) = \mu(B_{k+1})$ and $B_{k+1} = \varphi(B_k)$, where $B_k \subset M$ and $B_{k+1} \subset M$ are parameterized by smooth charts, $x_k(\xi) : [0,1] \to B_k$ and $x_{k+1}(\xi) : [0,1] \to B_{k+1}$, respectively.

explains Property (1), while the remaining four properties and their consequences are further explained and illustrated in the following sections.

As mentioned above, SRB measures are guaranteed to exist in Axiom A (or, uniformly hyperbolic) systems. Different rigorous studies indicate that uniform hyperbolicity is in fact not required for the existence of μ. For example, partially hyperbolic systems that have a mostly expanding [4] or contracting [25] central direction also admit SRB measures. In addition, many high-dimensional systems arising from the discretization of real-world PDE models behave as uniformly hyperbolic systems, per the *hyperbolicity hypothesis* [59].

In many engineering applications, the expected value of some physically relevant quantity $J \in L^1(\rho)$, i.e., $\int_M J \, d\mu$, is usually of interest. The major challenge in the field of sensitivity analysis is to find a parametric derivative of the expected value, which is critical in grid adaptation [96], optimization design [83] and uncertainty quantification [180]. Ruelle rigorously derived a closed-form expression for that derivative [148, 150],

$$\frac{d}{ds}\int_M J(x)\, d\mu(x) = \sum_{t=0}^{\infty} \int_M D(J \circ \varphi^t(x)) \cdot \chi(x)\, d\mu(x), \tag{3.48}$$

where χ denotes the derivative of φ with respect to the map parameter s[2], while D is the phase space differentiation operator. One could potentially apply a Monte Carlo algorithm to the integrals on the RHS of Eq. 3.48. However, owing to the butterfly effect, the direct evaluation of the integrand for a higher t is computationally infeasible. To illustrate this problem, let us consider the 2D Arnold's cat map $\varphi : [0,1]^2 \to [0,1]^2$ defined as

$$x_{k+1} = Ax_k \bmod 1, \quad A = \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix}, \tag{3.49}$$

and some smooth function $J(x)$. In Figure 3-13, we observe that even for a low t, $J \circ \varphi^t$ becomes highly-oscillatory, which implies that $\|D(J \circ \varphi^t)\|$ grows very fast ($\|\cdot\|$ denotes the Euclidean norm in $\mathbb{R}^n$). Due to the presence of positive Lyapunov exponents in chaotic systems, the rate of growth is in fact exponential. It means that Ruelle's formula is impractical for a direct Monte Carlo computation.

To circumvent this problem, one can split the perturbation vector χ and apply integration by parts to move the differentiation operator away from the composite function, as concisely described in [152]. Indeed, this approach gave rise to a form of the FDT theorem [149] widely used in statistical mechanics, and was also utilized to construct several numerical procedures approximating the linear response [32, 31, 127]. Based on the previous work, we carefully describe every step of the non-trivial partial integration process.

In case of integrals with respect to a non-uniform measure, integration by parts requires differentiating the measure itself. However, according to Property (3) and Property (4), $D\rho$ generally does not exist. In this section, let us assume χ equals a unit vector q that is tangent to the one-dimensional unstable manifold at every point on the manifold M. In a general case, $\chi \neq q$ and thus an extra step is required to regularize Ruelle's formula. This step involves a splitting of χ into two terms, such that one term belongs to unstable manifolds everywhere on M [152]. The reader is also referred to Chapter 5 for a detailed description of the entire process and relevant

[2]Note that χ is an n-dimensional analog of the perturbation f introduced in Chapter 2.

computational procedures. Thus, every integral from the RHS of Eq. 3.48 can be regularized through the following multi-step process (the description of each step follows the equation),

$$
\begin{aligned}
I &= \int_M DJ_t(x) \cdot q(x)\, d\mu(x) && (3.50)\\
&= \int_{M/U} \int_{U_x} DJ_t(s) \cdot q(s)\, d\tilde{\mu}_x(s)\, d\hat{\mu}(x) && (3.51)\\
&= \int_{M/U} \int_{U_x} \partial_q J_t(s)\, \tilde{\rho}_x(s)\, ds\, d\hat{\mu}(x) && (3.52)\\
&= \int_{M/U} \int_0^1 \partial_\xi J_t(s(\xi))\, \tilde{\rho}_x(s(\xi))\, d\xi\, d\hat{\mu}(x) && (3.53)\\
&= -\int_{M/U} \int_{U_x} J_t(s)\, \partial_q \tilde{\rho}_x(s)\, ds\, d\hat{\mu}(x) + (\text{boundary term}) && (3.54)\\
&= -\int_{M/U} \int_{U_x} J_t(s)\, \frac{\partial_q \tilde{\rho}_x}{\tilde{\rho}_x}(s)\, d\tilde{\mu}_x(s)\, d\hat{\mu}(x) + (\text{boundary term}) && (3.55)\\
&= -\int_M J_t(x)\, g(x)\, d\mu(x) + (\text{boundary term}), && (3.56)
\end{aligned}
$$

where $J_t(x) := J \circ (\varphi^t(x))$. To derive the final form of I, we perform the following steps. First (Step 3.51), we disintegrate μ on a measurable partition U determined by the geometry of unstable manifolds. The quotient measure $\hat{\mu}$ is defined such that for all Borel sets $B \subset M$,

$$\mu(B) = \int_{M/U} \tilde{\mu}_x(B \cap U_x)\, d\hat{\mu}(x),$$

where $\tilde{\mu}_x$ is a conditional SRB measure with density $\tilde{\rho}_x$. Subsequently, in Step 3.52, we use the measure-density relation, $d\tilde{\mu}_x = \tilde{\rho}_x\, ds$, where s denotes the path length as we move along U_x. In Step 3.53, we parameterize U_x, which gives rise to $ds = \|x'(\xi)\|\, d\xi$. Note the multiplicative factor is absorbed by the parametric derivative of J, because $\partial_\xi J = \|x'(\xi)\|\, \partial_s J$. Integration by parts is applied in Step 3.54, where the differentiation operator is moved from J to $\tilde{\rho}$. In Steps 3.55-3.56, we reshuffle terms and use the above identities again to simplify the final expression. Integration by parts also gives rise to a boundary term, which involves two integrals with respect to

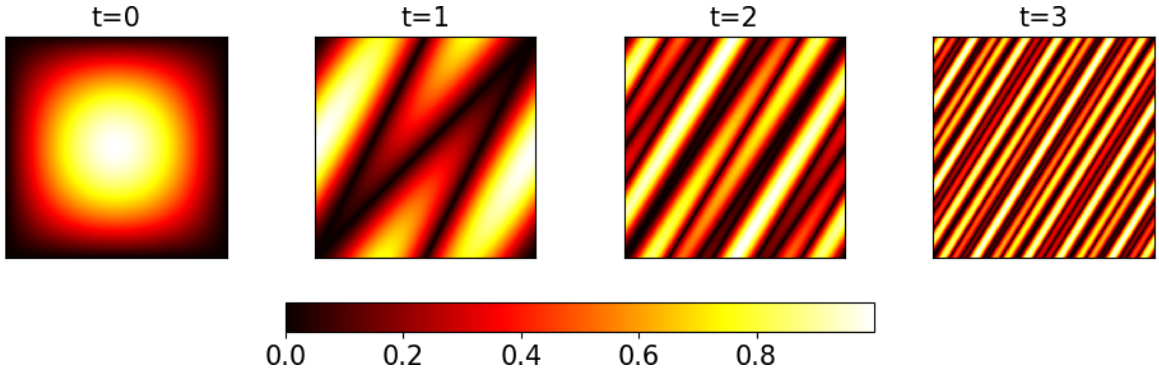

Figure 3-13: Evaluation of the composite function $J \circ \varphi^t(x)$ on the manifold $M = [0,1]^2$ at four consecutive steps t. In this case, the map φ is the Arnold's cat map (Eq. 3.49), while $J(x^{(1)}, x^{(2)}) = \sin(\pi x^{(1)}) \sin(\pi x^{(2)})$. This particular φ is a classical representative of an Anosov diffeomorphism.

the quotient measure of $J\,\tilde{\rho}_x$ evaluated at $\xi = 0$ and $\xi = 1$, respectively. From now on, we shall drop the subscript notation for conditional distributions; the tilde $\tilde{(\cdot)}$ notation shall imply the given distribution is restricted to a local unstable manifold. Note the boundary term,

$$(\text{boundary term}) = \int_{M/U} [\tilde{\rho}(\xi)\, J_k(\xi)]_{\xi=0}^{\xi=1}\, d\hat{\mu}(x), \tag{3.57}$$

can be expressed in terms of a regular volume integral over M of the divergence on unstable manifolds, which vanishes according to Theorem 3.1(b) of [148]. This is indeed a direct consequence of the fact the boundary terms across two neighboring rectangles of the Markov partition of M cancel out. To visualize this property, let us consider the Arnold's cat map (Eq. 3.49), for example. Despite its "artificial" discontinuities due to the modulo operator, this nonlinear transformation in fact maps a smooth torus to itself. One could arbitrarily change the boundaries of the square M in both phase space directions without modifying the map itself, and still describe the same torus.

Since

$$\int_M DJ_k \cdot q\, d\mu = -\int_M J_k\, g\, d\mu,$$

we can alternatively apply Monte Carlo to the RHS that involves the SRB density

gradient g [31, 159, 158],

$$g(x) = \frac{\partial_q \rho(x)}{\rho(x)} = \frac{\partial_q \tilde{\rho}(x)}{\tilde{\rho}(x)} = \partial_q \log \tilde{\rho}(x). \tag{3.58}$$

Note that the integrand appearing in the regularized version of I does not grow exponentially with t if J is bounded, which makes the sensitivity formula computable (immune to the butterfly effect). The integration by parts, as presented above, is generally useful if the integrand involves highly-oscillatory functions. The Monte Carlo integral example presented in [160] shows that the partial integration may reduce the number of samples by a few orders of magnitude to achieve the desired approximation error. Therefore, the computation of J might be beneficial not only in the context of Ruelle/S3/FDT-based methods for sensitivity approximation, but also in a general setting when the expected value of an ill-behaved quantity of interest in a chaotic system is needed. The following two sections focus on the computation of g for systems with an arbitrary number of positive LEs. The primary goal is to derive a recursive procedure compatible with Monte Carlo algorithms widely used in the field of chaotic dynamics.

3.7 Computing SRB density gradient for systems with one-dimensional unstable manifolds

In this section, we consider a generic n-dimensional, $n \in \mathbb{Z}^+$, uniformly hyperbolic dynamical system with one-dimensional unstable manifold governed by the C^2 diffeomorphic map $\varphi : M \to M$. M is thus a Riemannian manifold immersed in $\mathbb{R}^n$. There exists a measurable partition U of M such that each member of that partition, U_x, coincides with the unstable manifold that contains $x \in M$. In this particular case, each $U_x \subseteq M$ is geometrically represented by a curve embedded in $\mathbb{R}^n$. We strive to compute the directional derivative of the logarithmic SRB density g defined by Eq. 3.58.

3.7.1 Derivation of the iterative formula

The following notation is used throughout this section. Let $x_k(\xi) : [0,1] \to U_k \subset M$ denote a C^2 chart (diffeomorphic map) that describes the unstable manifold U_k, $k \in \mathbb{Z}$. For any k, the two charts: $x_k(\xi)$ and $x_{k+1}(\xi)$, defined respectively on U_k and U_{k+1}, are related as follows,

$$x_{k+1}(\xi) = \varphi(x_k(\xi)) \tag{3.59}$$

for all $\xi \in [0,1]$ (see Figure 3-12 for an illustration of an $n = 2$ case). We use $D\varphi$ and $D^2\varphi$ to respectively denote the Jacobian ($n \times n$ matrix) and Hessian ($n \times n \times n$ third-order tensor) of φ. Since φ is invertible, Eq. 3.59 can be viewed as a mathematical description of the evolution of SRB measure. For any observable f defined on M, evaluated along a certain trajectory, we use the following short-hand notation, $f \circ x_k(\xi) := f_k$. Derivatives of the chart with respect to the parameter ξ are denoted using the prime ($'$) symbol. The reference to the i-th component of an array (vector/matrix/tensor) is indicated inside round brackets located in the superscript; for example, $q^{(i)}$ denotes the i-th component of q. Finally, we use ∂_i to denote differentiation with respect to the i-th coordinate of phase space.

Let us parameterize U_k such that

$$\xi = \int_{\mathcal{C}_k(\xi)} \tilde{\rho}(x_k(\xi))\, ds, \tag{3.60}$$

where $\mathcal{C}_k(\xi)$ represents the segment of U_k between $x_k(0)$ and $x_k(\xi)$, which implies that $\mathcal{C}_k(1) \equiv U_k$. Consequently, $\tilde{\rho}_k$ is the conditional SRB density restricted to U_k satisfying $\tilde{\rho}_k = \rho_k / \int_{U_k} \rho_k\, ds$. We call it *measure-based* parameterization, as the value of the parameter ξ coincides with the value of the SRB measure at $x_k(\xi) \in U_k$. The variable transformation between ξ and the arc length s implies that

$$\tilde{\rho}(x_k(\xi))\, \|x_k'(\xi)\| = 1. \tag{3.61}$$

Note Eq. 3.61 is in fact a formula for the density change from a uniform to nonuniform distribution due to the nonlinear variable transformation $x_k(\xi)$. Since $\xi \in [0,1]$, $\tilde{\rho}_k$,

ξ, x_k can be respectively viewed as a probability density function (PDF), cumulative distribution (CDF), and inverse cumulative distribution (quantile function). Using the measure-based parameterization as described above, the SRB density gradient, defined in Eq. 3.58, can be expressed in terms of parametric derivatives of the chart, i.e.,

$$g(x_k(\xi)) = \frac{\partial_q \tilde{\rho}_k}{\tilde{\rho}_k}(x_k(\xi)) = -\frac{x_k'(\xi) \cdot x_k''(\xi)}{\|x_k'(\xi)\|^3}, \tag{3.62}$$

for any $\xi \in [0,1]$. Here, the derivative ∂_q is computed in the direction of increasing value of ξ. The reader is referred to Section 3.2.3, where Eq. 3.62 is derived by differentiating Eq. 3.61, and comprehensively described using various numerical examples.

We notice $x_k'(\xi) = \|x_k'(\xi)\|\, q(x_k(\xi))$, and rewrite Eq. 3.62 to

$$g(x_k(\xi)) = -q(x_k(\xi)) \cdot \frac{x_k''(\xi)}{\|x_k'(\xi)\|^2} := -q(x_k(\xi)) \cdot a(x_k(\xi)) = -q_k \cdot a_k. \tag{3.63}$$

Eq. 3.63 indicates that the magnitude of the SRB density gradient equals the length of the projection of the (re-scaled) curve acceleration vector on the line tangent to the curve. We now use Eq. 3.59, differentiate it twice with respect to ξ, and apply the chain rule to obtain the following expression,

$$x_{k+1}''(\xi) = D^2\varphi(x_k(\xi))(x_k'(\xi), x_k'(\xi)) + D\varphi(x_k(\xi))\, x_k''(\xi), \tag{3.64}$$

which means that

$$\begin{aligned} a(x_{k+1}(\xi)) =& \frac{x_{k+1}''(\xi)}{\|x_{k+1}'(\xi)\|^2} = \\ & \frac{\|x_k'(\xi)\|^2 D^2\varphi(x_k(\xi))\,(q(x_k(\xi)), q(x_k(\xi)))}{\|x_{k+1}'(\xi)\|^2} + \frac{D\varphi(x_k(\xi))x_k''(\xi)}{\|x_{k+1}'(\xi)\|^2}. \end{aligned} \tag{3.65}$$

The bilinear form that appears in the first term on the RHS of Eq. 3.65 can be expressed using Einstein's summation convention, i.e.,

$$[D^2\varphi(q,q)]^{(ijk)} = \partial_i\partial_j\varphi^{(k)}\, q^{(i)}\, q^{(j)}.$$

Given

$$\|x'_{k+1}(\xi)\| = r(x_k(\xi))\|x'_k(\xi)\|$$

and

$$r(x_k(\xi)) = \|D\varphi(x_k)\ q(x_k(\xi))\|,$$

we conclude that

$$a_{k+1} = \frac{(D^2\varphi)_k(q_k, q_k) + (D\varphi)_k\ a_k}{r_k^2}. \tag{3.66}$$

From the parametric derivative of Eq. 3.59 and the definition of $r(x_k(\xi))$, the recursion

$$q_{k+1} = \frac{D\varphi_k\ q_k}{r_k} \tag{3.67}$$

automatically follows. We emphasize the fact the above procedure for g (involving Eq. 3.63, 3.66, 3.67) is completely analogous to the algorithm proposed in Section 3.4.2, which was meant for simple Lebesgue measures evolving due to a generic non-chaotic diffeomorphisms. Here, however, we consider the evolution of the SRB measure in a chaotic system. Due to the butterfly effect, the tangent solution exponentially increases in norm. Therefore, we need the normalizing factor r in the iterative formula for a and q along the trajectory. Since φ is uniformly hyperbolic, the solution to the tangent equation in Eq. 3.67 converges exponentially in k to the backward Lyapunov vector that is tangent to the unstable manifold regardless of the choice of an initial condition q_0. Under the same assumption, the recursion in Eq. 3.66 for the acceleration vector a also uniformly converges to the true solution at an exponential rate for any initial condition a_0 bounded in norm. The reader is referred to Lemma 7.7 in [31] for the proof of the preceding statement.

To summarize, using the measure-based manifold parameterization, we derived a simple recursive procedure for the SRB density gradient that exponentially converges in case of uniformly hyperbolic systems and does not depend on initial conditions. As for now, we restrict ourselves to systems with one-dimensional unstable manifolds. Our main intention here is to introduce basic concepts before we move to general cases in Section 3.8.

3.7.2 Numerical example: computing SRB density gradient on straight unstable manifolds

As a pedagogical example, let us consider a family of n-dimensional maps, $n \in \mathbb{Z}^+$, whose unstable manifolds are straight and, without loss of generality, aligned with the first coordinate of the phase space. Certainly, this family includes, but is not limited to, all one-dimensional chaotic maps. In this particular case, $q^{(i)} = \delta^{(i1)}$, where δ denotes the Kronecker delta. Consequently, the parametric derivative of the chart $x_k(\xi)$, for any k, has all zero entries except the first one and, therefore, $r(x_k(\xi)) = |\partial_1 \varphi^{(1)}(x_k(\xi))|$. Thus, our recursive algorithm for g, which involves Eq. 3.63, 3.66, and 3.67, reduces to a single scalar iterative formula,

$$g(x_{k+1}(\xi)) = \frac{g(x_k(\xi))}{\partial_1 \varphi^{(1)}(x_k(\xi))} - \frac{\partial_1^2 \varphi^{(1)}(x_k(\xi))}{(\partial_1 \varphi^{(1)}(x_k(\xi)))^2} \tag{3.68}$$

for all $\xi \in [0, 1]$. We were allowed to drop the absolute values, because $x_k'(\xi) > 0$, which is a consequence of our choice of the manifold parameterization. In this simple event of a straight unstable manifold, only two scalars are required to advance the iteration, i.e., first- and second-order derivative (in phase space) of the first component of φ, since the map is expanding only in one direction. This result is fully consistent with early non-systematic attempts to construct such a procedure for g in [158, 159] (see Section 2.6.3). The previous studies used the measure preservation property to derive an expression analogous to Eq. 3.68.

To verify the correctness of our procedure, we consider the 2D perturbed Baker's map $\varphi : M \to M$, with $M = [0, 2\pi]^2$, defined as follows [31],

$$\begin{aligned} x_{k+1} = \varphi(x_k) = \Bigg(& \begin{bmatrix} 2x_k^{(1)} \\ x_k^{(2)}/2 + \pi \lfloor x_k/\pi \rfloor \end{bmatrix} + \\ & \begin{bmatrix} s_1/2 \, \sin(x_k^{(1)}/2) + s_2/2 \, \sin(2x_k^{(1)}) \, \sin(x_k^{(2)}) \\ s_3 \, \sin(x_k^{(2)}) + s_4/2 \, \sin(2x_k^{(1)}) \, \sin(x_k^{(2)}) \end{bmatrix} \Bigg) \bmod 2\pi, \end{aligned} \tag{3.69}$$

where s_1, s_2, s_3, s_4 are real-valued map parameters. If all of them are zero, we obtain

the classical Baker's map (first term of the RHS of Eq. 3.69), which is named after the kneading operation that bakers apply to a two-dimensional square dough. In particular, the dough is first stretched horizontally (in the unstable direction) by a constant factor, then compressed vertically (in the stable direction) by the same factor, and so forth. The square-shaped domain is stretched to a 2×1 rectangle, cut into two squares, which are subsequently stacked horizontally. The Baker's map is an invertible chaotic map with one positive and one negative Lyapunov exponent.

By introducing an extra term proportional to the four parameters, we perturb the kneading operation in the direction not necessarily aligned with the phase space directions. Indeed, by manipulating these parameters' values, we can control the shape of the unstable manifold, which gives us an excellent study case in the context of the SRB gradient computation. Notice, for example, if $s_4 = 0$ and s_3 is sufficiently small, the iteration in Eq. 3.67 produces q_k whose second coordinate, $q_k^{(2)}$, converges exponentially to zero with k. In this case, therefore, unstable manifolds are straight and aligned with the $x^{(1)}$-axis. We use this observation to design our first numerical test.

In the first experiment, we consider the Baker's map defined by Eq. 3.69 with $s_1 = s_3 = s_4 = 0$ and $s_2 = 0.4$. The left-hand side plot in Figure 3-14 illustrates the normalized SRB distribution corresponding to this parameter choice, which represents the probability of the trajectory passing through each square bin everywhere on M (see the caption of Figure 3-14 for more details; for completeness, we also included a case with $s_4 \neq 0$). We observe a smooth behavior of the SRB distribution with respect to $x^{(1)}$ at any vertical level $x^{(2)}$. However, as we travel vertically, in the stable direction, the SRB distribution varies sharply. These radically different behaviors are typical symptoms of Property (3) and Property (4) of μ described in Section 3.6, and they can also be observed in Figure 3-15, where the conditional and marginal SRB distributions are plotted, using data from Figure 3-14.

In Figure 3-15, we also plot the SRB density gradients defined on five different unstable manifolds. To compute g, the simplified recursion from Eq. 3.68 was directly applied. To validate our computation, we approximated g by applying the central

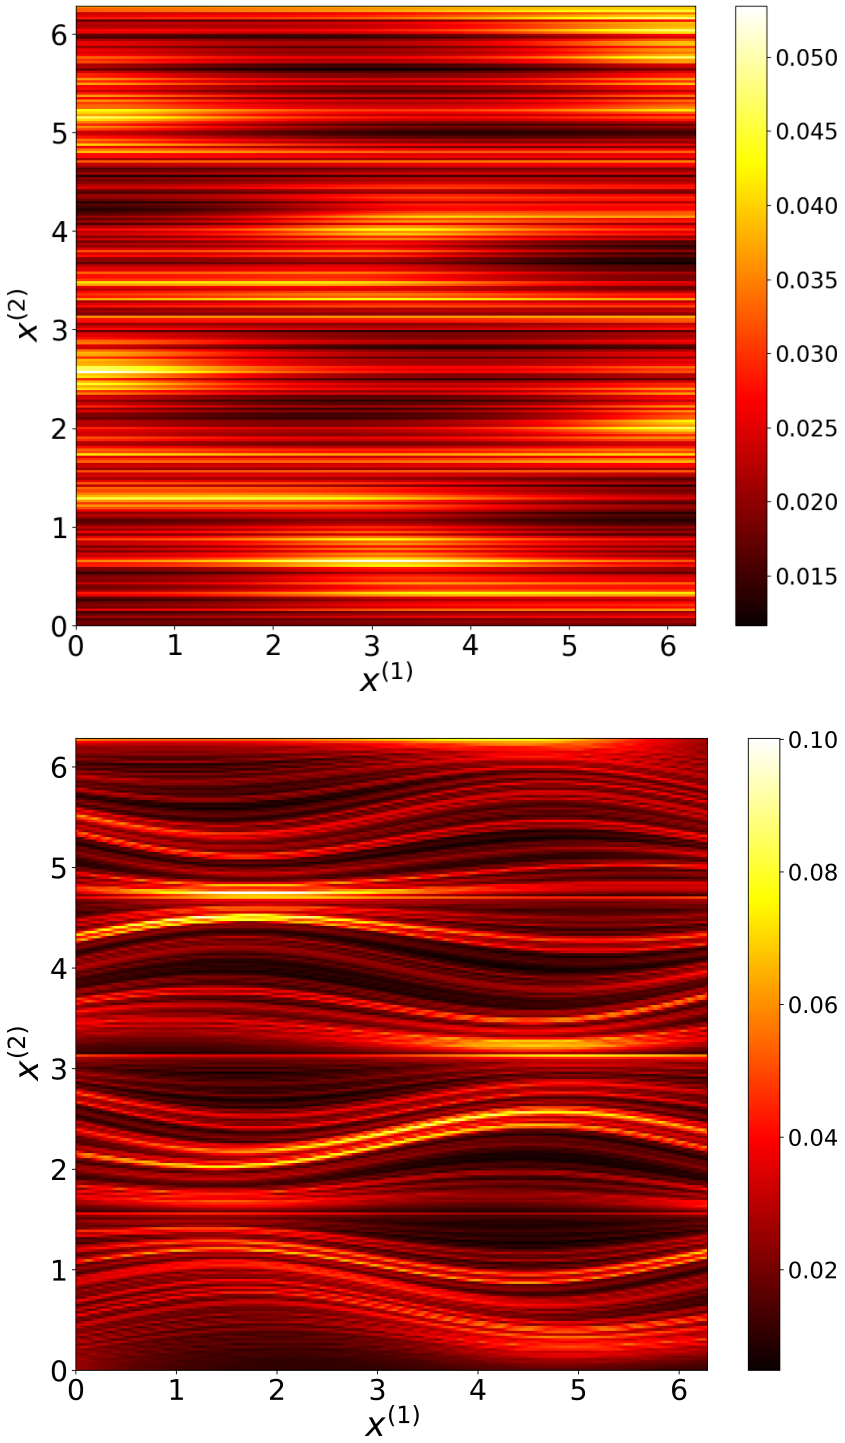

Figure 3-14: SRB distribution of the Baker's map with $s_1 = s_3 = s_4 = 0$, $s_2 = 0.4$ (left plot) and $s_1 = s_2 = s_3 = 0$, $s_4 = 0.4$ (right plot). We divided M into 256^2 rectangular bins of equal width and counted the number of times the trajectory passed through each of these bins. In this experiment, we generated 8000 trajectories of length $209,715,200$, which gives us the total of approximately $1.68 \cdot 10^{12}$ samples.

finite-difference method to SRB densities plotted above. We observe a good agreement between the results obtained with these two different approaches, which confirms the correctness of our algorithm.

To conclude, in case of straight unstable manifolds, the SRB density gradient can be computed using the simplified recursive relation along trajectory (Eq. 3.68), which we verify through finite-differencing. This iteration is computationally cheap, as it involves solving a scalar tangent equation featuring both the first and second derivative of the first component of φ. In Appendix 3.10.1, we show Eq. 3.68 can also be applied to popular one-dimensional maps that are non-injective. We argue that certain non-measure-preserving transformations have their higher-dimensional analogs similar to the classical Baker's map. Appendix 3.10.2 presents a numerical study confirming the hyperbolicity of the Baker's map.

3.8 Computing SRB density gradient for systems with general unstable manifolds

We shall generalize the concepts introduced in Section 3.7 to systems with m-dimensional unstable manifolds, $m \in \mathbb{Z}^+$. In other words, we consider general n-dimensional chaotic systems that have m positive LEs, $1 \leq m \leq n$. In this setting, the chart $x_k(\xi)$, $k \in \mathbb{Z}^+$, is a diffeomorphism that maps an m-dimensional hypercube, $[0,1]^m$, to the local unstable manifold $U_k \subset M$. For example, if $m = 2$ and $n = 3$, then the system has two positive LEs and its unstable manifolds are surfaces immersed in $\mathbb{R}^3 \supset M$.

3.8.1 Derivation of the iterative formula

As introduced above, let us consider an m-dimensional smooth unstable manifold U_k described by the chart $x_k(\xi) : [0,1]^m \to U_k \subset M$. The vectors $x_k = [x_k^{(1)}, ..., x_k^{(n)}]^T$ and $\xi = [\xi^{(1)}, ..., \xi^{(m)}]^T$ have n and m components, respectively, and $0 \leq \xi^{(i)} \leq 1$, $i = 1, ..., m$. We use $\nabla_\xi x_k(\xi)$ to denote the parametric gradient tensor of the chart.

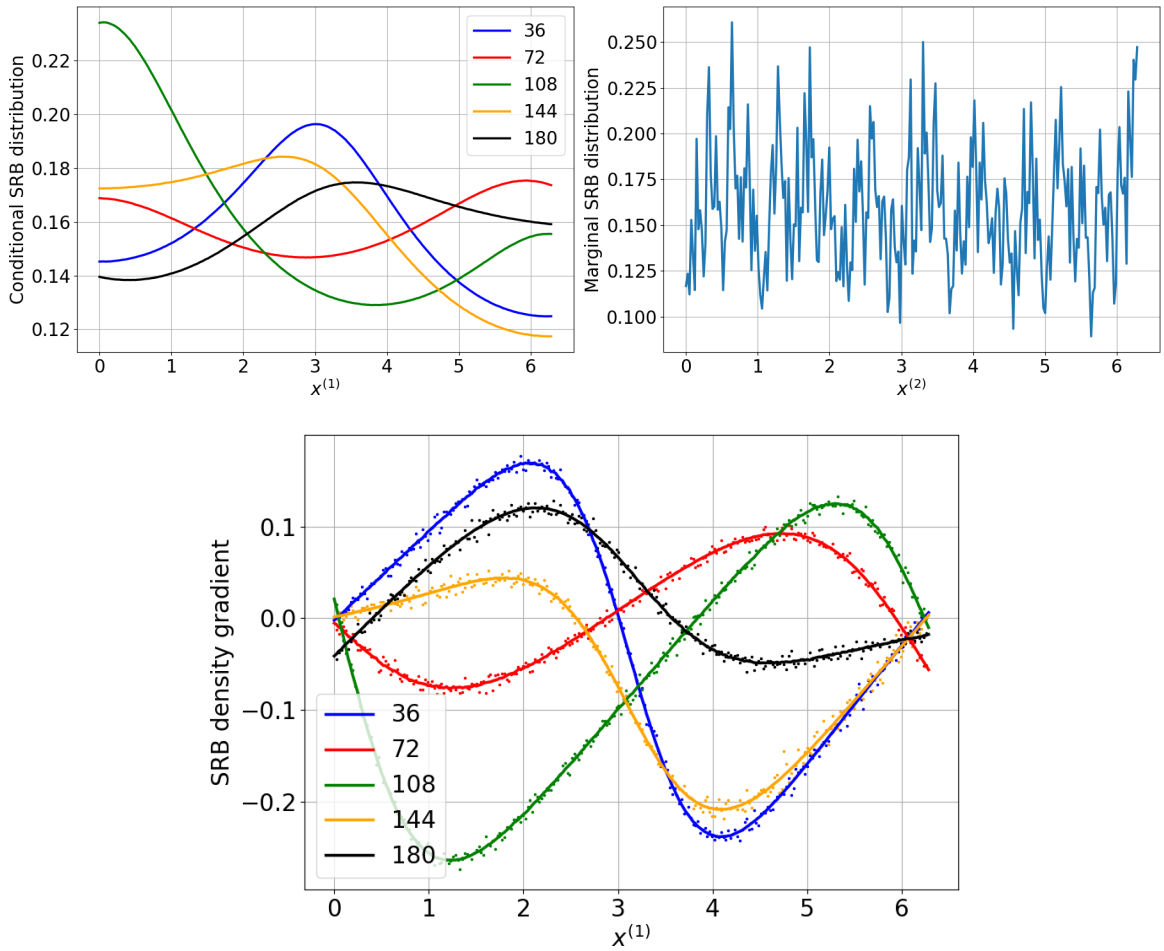

Figure 3-15: Upper left plot: conditional SRB distributions (SRB densities) corresponding to five different unstable manifolds. The numbers $36, 72, 108, 144, 180$ appearing in the legend represent the index of the horizontal bin row. For example, the red line corresponds to the SRB density defined on the unstable manifold at $x^{(2)} \approx 72/256 \cdot 2\pi \approx 1.76$. Upper right plot: marginal SRB distribution obtained through integrating the first coordinate out. Lower plot: SRB density gradient g corresponding to SRB densities plotted in Figure 3-15. The g function was computed using two distinct approaches: through the simplified trajectory-based recursion (Eq. 3.68)(solid lines), and the central finite-difference method (dots). The oscillation of the finite-difference approximation is a manifestation of the statistical noise.

The i-th column of $\nabla_\xi x_k(\xi)$ contains the derivative of $x_k(\xi)$ with respect to $\xi^{(i)}$, i.e., $\partial_{\xi^{(i)}} x_k(\xi)$. For any Borel subset $V \subset [0,1]^m$ such that $x_k(V) = B_k \subset U_k$, the SRB measure-density relation can be expressed as follows,

$$\mu(V) = \int_{B_k} \tilde{\rho}_k(x)\, d\omega(x), \tag{3.70}$$

where $d\omega(x)$ denotes the natural volume element defined everywhere on U_k. Analogously to the 1D case described in Section 3.7, $\tilde{\rho}_k$ represents the conditional SRB density defined on U_k. If we QR-factorize the parametric gradient of $x_k(\xi)$,

$$\nabla_\xi x_k(\xi) = Q(x_k(\xi))\; R(x_k(\xi)) \tag{3.71}$$

at any $\xi \in [0,1]^m$, the density conservation property could be expressed as

$$\tilde{\rho}(x_k(\xi))\; |\det R(x_k(\xi))| = 1, \tag{3.72}$$

which is a generalization of Eq. 3.61. Furthermore, by differentiating Eq. 3.72 with respect to ξ and applying a non-trivial chain rule, we obtain

$$\begin{aligned} g^{(i)}(x_k(\xi)) :=& \partial_{s^{(i)}(x_k(\xi))} \log \tilde{\rho}(x_k(\xi)) = \\ & - \frac{\operatorname{tr}\left(Q^T(x_k(\xi))\; \partial_{\xi^{(i)}} \nabla_\xi x_k(\xi)\; R^{-1}(x_k(\xi))\right)}{\|\partial_{\xi^{(i)}} x_k(\xi)\|}, \end{aligned} \tag{3.73}$$

or, equivalently,

$$g^{(i)}(x_k(\xi)) = -\frac{Q^{(:j)}(x_k(\xi)) \cdot \partial_{\xi^{(i)}} \partial_{\xi^{(k)}} x_k(\xi)\; (R^{-1})^{(kj)}(x(\xi))}{\|\partial_{\xi^{(i)}} x_k(\xi)\|} \tag{3.74}$$

for all $\xi \in [0,1]^m$, where the repeated indices imply summation (Einstein's convention), while the superscript $(:i)$ denotes the i-th column of a matrix. This expression was obtained by employing the orthogonality of Q and upper-triangular structure of R. It is computationally convenient as it does not involve parametric derivatives of the determinant of R. Note that Eq. 3.73 represents a formula for log-density

derivatives in the directions indicated by vectors tangent to parametric isolines. In particular, $s^{(i)}$ represents the tangent vector of the isoline (curve) associated with $\xi^{(i)}$, $i = 1, ..., m$. The reader is referred to Section 3.3.1 for a step-by-step derivation of Eq. 3.72-3.74. We acknowledge that a similar "trace" formula was derived in the context of computing the unstable divergence in [127].

The purpose of this section is to derive an iterative (trajectory-driven) procedure for g. Analogously to the derivation in Section 3.7, we combine Eq. 3.32, the evolution equation

$$x_{k+1}(\xi) = \varphi(x_k(\xi)), \tag{3.75}$$

and apply the chain rule. The 1D case ($m = 1$) was computationally simpler because the tangent equations for a and q were regularized by the scalar r every time step preventing the tangent solutions from blow-ups due to the positive LE . Here, we need to compute all first- and second-order parametric derivatives of the chart to compute g. Thus, we aim to regularize tangent equations in a fashion analogous to the approach in Section 3.7. To achieve this goal without violating the consistency of the scheme, we recursively orthonormalize the parametric gradient through an iterative linear transformation of the parameterization and fixing $\xi = 0$. In particular, we change variables from step k to $k + 1$ such that

$$\xi_{k+1} = R_{k+1}(x_{k+1}(\xi_k))\, \xi_k.$$

Note that at $\xi = 0$ we stay on the same trajectory despite the transformation. This particular choice of ξ does not restrict our algorithm to concrete trajectories. Indeed, we want to "visit" all infinitesimally small μ-typical regions of the attractor after an infinite number of time steps, regardless of the choice of the initial condition. Therefore, we can always linearly re-scale the feasible space of ξ such that $\xi = 0$ for our arbitrary choice of the initial condition. To simplify the notation, we skip the argument whenever $\xi = 0$; for example, we use the short-hand notation $x_k(0) := x_k$, $Q(x_k(0)) := Q_k$, and so forth. Thanks to this particular coordinate transformation, the parametric gradient is automatically orthonormalized, because the chain rule

implies that

$$\nabla_{\xi_{k+1}} x_{k+1} = \nabla_{\xi_k} x_{k+1}\, R_{k+1}^{-1} = Q_{k+1}, \tag{3.76}$$

or, equivalently,

$$\partial_{\xi_{k+1}^{(i)}} x_{k+1} = \partial_{\xi_k^{(j)}} x_{k+1}\, (R_{k+1}^{-1})^{(ji)} = Q_{k+1}^{(:i)}. \tag{3.77}$$

It means that the parametric gradient of the chart has an orthonormal basis of the column space in the updated coordinate system. Note the R matrix represents the Jacobian of the step-to-step parametric transformation, i.e.,

$$R_{k+1} = \partial \xi_{k+1} / \partial \xi_k.$$

Through recursive re-orthonormalization of the coordinate chart, we derive a consistent and recursive scheme for the unstable basis. Assuming the chart gradient is orthogonal at step k, we take the gradient of Eq. 3.75 with respect to the ξ_k (i.e., the coordinates that guarantee $\nabla_{\xi_k} x_k = Q_k$) to conclude that

$$Q_{k+1}\, R_{k+1} = \nabla_{\xi_k} x_{k+1} = \nabla_{\xi_k} \varphi(x_k) = D\varphi_k \nabla_{\xi_k} x_k = D\varphi_k Q_k, \tag{3.78}$$

which further implies that

$$Q_k = D\varphi_{k-1}\, D\varphi_{k-2} \ldots D\varphi_0\, Q_0\, R_1^{-1} \ldots R_{k-1}^{-1}\, R_k^{-1}. \tag{3.79}$$

Notice that Eq. 3.78 is obtained by differentiating the evolution equation first, and then applying the orthonormalized coordinates. Eq. 3.79 coincides with the already known algorithm for Gram-Schmidt (backward) Lyapunov vectors [90, 65], which converges exponentially fast in uniformly hyperbolic systems as $k \to \infty$ for any bounded Q_0. For a sufficiently large k, the set of column vectors of Q_k approximates an orthonormal basis of the unstable (expanding) subspace E_k^u of the tangent space TM_k. Specific directions of backward Lyapunov vectors at time step k, however, depend on the choice of Q_0. Therefore, in this case, the "convergence" should be understood that, for any orthonormal Q_0, the column space of Q_k is guaranteed to

coincide with some orthonormal basis of E_k^u if $k \to \infty$. A similar procedure can be used to compute all n Lyapunov vectors, including those corresponding to the negative LEs, spanning the stable (contracting) subspace E_k^s. In uniformly hyperbolic systems,

$$TM_k = E_k^u \oplus E_k^s$$

at every k, and both the subspaces are $D\varphi$-invariant (or *covariant*). The covariance property implies that the product $D\varphi_k Q_k$ outputs m vectors that belong to the unstable subspace of the tangent space at the next time step, TM_{k+1}. In general, the new vectors are not orthonormal. By performing the QR factorization, however, we obtain an orthonormal basis of the unstable subspace at $k+1$. The components of R_{k+1} contain projections of the column vectors of $D\varphi_k Q_k$ onto the basis vectors of E_{k+1}^u.

There are two other important consequences of the recursive coordinate change. First, notice that the tangent lines of parametric isolines are now orthogonal. Therefore,

$$g^{(i)} = \partial_{s^{(i)}} \log \tilde{\rho}_k = \partial_{Q^{(:i)}} \log \tilde{\rho}, \tag{3.80}$$

where $Q^{(:i)} := q^i$ denotes the i-th column of Q. The second consequence is a dramatic simplification of Eq. 3.32. Indeed, in the orthonormalized system, the R matrix reduces to the identity matrix, while the norm of each column of the parametric gradient equals 1. This gives rise to the following expression for g,

$$\begin{aligned} g_{k+1}^{(i)} = &- \operatorname{tr}\left(Q_{k+1}^T \, \partial_{\xi_k^{(i)}} \nabla_{\xi_k} x_{k+1}\right) = \\ &- Q_{k+1}^{(:j)} \cdot \partial_{\xi_{k+1}^{(i)}} \partial_{\xi_{k+1}^{(j)}} x_{k+1} := -Q_{k+1}^{(:j)} \cdot a_{k+1}^{(i,j)}. \end{aligned} \tag{3.81}$$

To complete the algorithm, we also need to recursively compute the Hessian of the coordinate cart, i.e., all possible parametric second-order derivatives of $x(\xi)$. We usually refer to those vectors as *acceleration vectors* or *curvature vectors*. They can

be computed recursively by differentiating Eq. 3.75 twice, which gives rise to

$$\begin{aligned}\partial_{\xi_k^{(i)}}\partial_{\xi_k^{(j)}} x_{k+1} &= D^2\varphi_k(\partial_{\xi_k^{(i)}} x_k, \partial_{\xi_k^{(j)}} x_k) + D\varphi_k\, \partial_{\xi_k^{(i)}}\partial_{\xi_k^{(j)}} x_k \\ &= D^2\varphi_k(Q_k^{(:i)}, Q_k^{(:j)}) + D\varphi_k\, a_k^{(i,j)}.\end{aligned} \tag{3.82}$$

Again, we highlight the fact that Eq. 3.82 was obtained by first differentiating the original evolution equation (first line), and then enforcing the orthonormality of the chart gradient (second line). To maintain the consistency of the entire scheme, we need to change the differentiation coordinates on the left-hand side of Eq. 3.82. Consequently, we apply the chain rule to derive a relation for the Hessian of x_{k+1}, represented by an $n \times m \times m$ tensor, in the updated coordinate system:

$$\partial_{\xi_{k+1}^{(i)}}\partial_{\xi_{k+1}^{(j)}} x_{k+1} = \partial_{\xi_k^{(p)}}\partial_{\xi_k^{(q)}} x_{k+1}\ (R_{k+1}^{-1})^{(pi)}\ (R_{k+1}^{-1})^{(qj)}. \tag{3.83}$$

Note that in order to compute the SRB density gradient at step $k+1$, we need to apply the Hessian re-scaling described by Eq. 3.83 to retrieve a_{k+1}. We summarize this algorithm and carefully analyse its computational properties in Section 3.8.2.

We observe the general recursion for the acceleration vector a in Eq. 3.82 can be simplified to its one-dimensional counterpart in Eq. 3.66 if $m = 1$. In this chapter, we resort to an empirical study of the convergence of our algorithm (see Section 3.8.2). In Chapter 5, we rigorously prove the stability and convergence of the re-scaled iteration combining Eq. 3.83 and Eq. 3.82. Similar analysis of systems with one-dimensional unstable manifolds was presented in [31, 33].

We reiterate the fact that the recursive change of coordinates has a "stabilizing" effect. Note that both tangent iterations for the Lyapunov bases and acceleration vectors respectively apply Jacobians and Hessians every time step leading to exponential increase in norm of the corresponding products. On the other hand, these exploding products are counterbalanced by R^{-1} matrices, which are by-products of the chain rule. Therefore, this re-scaling operation does not only guarantee the consistency of the scheme, but also eliminates the butterfly effect from the underlying recursions. In Chapter 5, we provide rigorous arguments on the numerical stability.

3.8.2 General algorithm for systems with multi-dimensional unstable manifolds

We provide a practicable algorithm based on the derivation presented in the previous section. In addition, we carefully analyse its computational cost, memory requirements, and numerically investigate its convergence. Algorithm 1 summarizes all steps necessary to numerically compute the SRB density gradient at N points along a trajectory initiated at $x_0 \in M \subset \mathbb{R}^n$. The only optional step is included in Line 1; this step is meant to compute the dimension of the unstable subspace/manifold m. For many chaotic maps, this parameter is known *a priori* and therefore Line 1 can be skipped. If this is not the case, however, one can apply Benettin et al.'s numerical procedure [15] to approximate a subset of the spectrum of Lyapunov exponents. This procedure requires solving $i \in \mathbb{Z}^+$ homogeneous tangent equations to identity i largest LEs. The parameter T represents the trajectory length and affects the accuracy of LE approximation. If the LE spectrum is evidently separated from the origin (i.e., the value of 0), then T does not need to be large. Lines 3-22 of Algorithm 1 represent the main time for-loop that computes the g vector at one point on the manifold per iteration. Inside this loop, we distinguish five major stages: 1) advancing first-order tangent equation and QR factorization (Eq. 3.79), 2) advancing second-order tangent equations (Eq. 3.82), 3) inverting the R matrix and re-scaling the acceleration vector a (Eq. 3.83), 4) evaluating g (Eq. 3.81), and 5) transitioning to the next time step; updating the Jacobian and Hessian.

Table 22 summarizes the computational cost of Algorithm 1. The third column of this table includes the number of the floating point operations required in each stage as a function of the trajectory length (N or T), system dimension n, and unstable manifold dimension m. Note the third column includes only the leading term of the flop count. The final two stages involve evaluations of nonlinear equations defined by φ and thus their computational cost is problem-dependent. In many physics-inspired chaotic systems, the cost of Lines 20-21 is relatively low. Consider the Lorenz '63 system discretized using the Euler scheme, for example. In this case, we can think

Algorithm 1: SRB density gradient

Input : N, T, x_0, $n = \text{size}(x_0)$

1 $m = \text{Benettin}(T)$ **if** m unknown;

2 Randomly generate Q_0, $a_0^{(i,j)}$ such that $\text{ncol}(Q_0) = m$, $\text{nrow}(Q_0) = \text{size}(a_0^{(i,j)}) = n$, $Q_0^T Q_0 = I$, and $i, j = 1, ..., m$;

3 **for** $k = 0, ..., N-1$ **do** `// main time loop`

4 $\quad S_k = D\varphi_k \, Q_k$;

5 $\quad$ QR-factorize: $Q_{k+1} \, R_{k+1} = S_k$;

6 $\quad$ Invert R_{k+1};

7 $\quad$ **for** $i = 1, ..., m$ **do** `// 2nd-order tangent equations`

8 $\quad\quad$ **for** $j = 1, ..., i$ **do**

9 $\quad\quad\quad \tilde{a}_{k+1}^{(i,j)} = D^2\varphi_k(Q_k^{(:i)}, Q_k^{(:j)}) + D\varphi_k \, a_k^{(i,j)}$;

10 $\quad\quad$ **end**

11 $\quad$ **end**

12 $\quad$ **for** $i = 1, ..., m$ **do** `// re-scaling`

13 $\quad\quad$ **for** $j = 1, ..., i$ **do**

14 $\quad\quad\quad a_{k+1}^{(i,j)} = \tilde{a}_{k+1}^{(p,q)} \, (R^{-1})_{k+1}^{(pi)} \, (R^{-1})_{k+1}^{(qj)}$;

15 $\quad\quad$ **end**

16 $\quad$ **end**

17 $\quad$ **for** $i = 1, ..., m$ **do** `// evaluating` g

18 $\quad\quad g_{k+1}^{(i)} = -Q_{k+1}^{(:j)} \cdot a_{k+1}^{(i,j)}$;

19 $\quad$ **end**

20 $\quad x_{k+1} = \varphi(x_k)$;

21 $\quad$ Evaluate: $D\varphi_{k+1}$ and $D^2\varphi_{k+1}$;

22 **end**

Output: $g_k^{(i)}$, $i = 1, ..., m$, $k = 1, ..., N-1$

of φ as a time discretization operator of the continuous system. The Jacobian $D\varphi$ involves a collection of linear terms proportional to the coordinates of x, while the Hessian $D^2\varphi$ is constant. In many scientific/engineering applications, PDE models are discretized in space using schemes with local support (such as the finite element method), which implies the resulting Jacobians and Hessians of the fully-discretized system are sparse. Therefore, in these special cases, the cost of the most expensive stage of Algorithm 1, which involves second-order tangent equations, can potentially be reduced to $N \, n \, m^2$. Table 22, however, reflects the worst-case scenario in which no sparsity patterns occur. We also highlight the fact that in many high-dimensional chaotic systems $m \ll n$ [18]. Thus, if n is large, the re-scaling stage (Lines 12-16) is

Table 3.1: Computational cost of Algorithm 1.

Stage Name	Line No.	Total Cost
Computing m (Benettin's algorithm)	1	$T\,n^2 m$
Generating initial conditions	2	–
Advancing first-order tangent equations	4	$N\,n^2\,m$
QR factorization (Householder)	5	$N\,n\,m^2$
Inverting R	6	$N\,m^3$
Advancing second-order tangent equations	7-11	$N\,n^3\,m^2$
Re-scaling a	12-16	$N\,n\,m^3$
Computing g	17-19	$N\,n\,m^2$
Advancing primal equation	20	Varies
Evaluating Jacobian and Hessian	21	Varies

rather cheaper than the second-order tangent equation stage (Lines 7-11).

We conclude that the leading term of the total flop count of Algorithm 1 is proportional to $N\,n^3\,m^2$ in a general chaotic system. In many real-world problems, however, the final cost can be significantly reduced if one takes the advantage of the system's special structure. Our algorithm is moderately cheap in terms of the memory requirements. The most memory-consuming structure is the Hessian which, in the worst-case scenario, requires storing n^3 floats. As we pointed out above, however, in practical high-dimensional models, the actual "size" of the Hessian might be dramatically smaller. Note also that, in order to advance tangent equations, we need to store m n-dimensional basis vectors (i.e., column vectors of Q) and $\sim 1/2\,m^2$ acceleration vectors. The $1/2$ factor is a consequence of the assumed smoothness of the coordinate chart, which implies $a^{(i,j)} = a^{(j,i)}$ everywhere on the manifold. Notice also that our procedure is in fact a one-step method, which means that all quantities at step $k+1$ require data only from step k. We do not need to store data generated at previous time steps.

Finally, we perform a numerical test to investigate the convergence properties of Algorithm 1. For this purpose, we use the Baker's map introduced in Eq. 3.69, as

well as its 3D version $\varphi : [0, 2\pi]^3 \to [0, 2\pi]^3$ defined as

$$x_{k+1} = \varphi(x_k) = \left(\begin{bmatrix} 2\, x_k^{(1)} \\ 3\, x_k^{(2)} \\ x_k^{(3)}/6 + \pi \lfloor x_k^{(1)}/\pi \rfloor + \pi/3 \lfloor x_k^{(2)}/(2\pi/3) \rfloor \end{bmatrix} + \begin{bmatrix} s_1 \sin(2x_k^{(1)}) \sin(3/2\, x_k^{(2)}) \\ s_2 \sin(x_k^{(1)}) \sin(3\, x_k^{(2)}) \\ s_3 \sin(6\, x_k^{(3)}) \end{bmatrix} \right) \bmod 2\pi, \tag{3.84}$$

which we shall refer to as the 3D Baker's map. The three scalars s_1, s_2, s_3 are real-valued input parameters of that map. This is an invertible chaotic map with two positive and one negative LEs, and seemingly hyperbolic behavior (see Appendix 3.10.2 for more details). This map has two expanding directions, along the $x^{(1)}$ and $x^{(2)}$ axes, and one contracting direction along the third axis. Analogously to its 2D counterpart, this map models the kneading operation. The dough is extended by the factor 2 and 3 along the two orthogonal directions on the $x^{(1)} - x^{(2)}$ plane, cut into $2 \cdot 3 = 6$ squares, which are subsequently stacked in the order defined by the floor functions. These history-dependent floor functions are used to guarantee the invertibility of the nonlinear transformation by periodically distributing the third component of x_{k+1} across $[0, 2\pi]$. Higher-dimensional Baker's maps have been widely used in image encryption as a convenient generalization of Bernoulli shifts [138, 118].

To analyze the convergence, we generate three sufficiently long trajectories started at randomly chosen initial conditions x_0. For each of these trajectories, we run two independent simulations with different, randomly chosen initial conditions for the tangent equations (see Line 2 of Algorithm 1). Motivated by the rigorous studies, we investigate if (and how) the difference between the SRB density gradients computed along a single trajectory but using different initial conditions for tangent equations decreases in norm as we advance the iteration. In particular, we compute $\|g_{k,1} - g_{k,2}\|$, $k = 0, 1, 2, ...$ for two random initial condition choices for tangent equations per trajectory, labelled as 1 and 2. The relation between that norm and time step k for

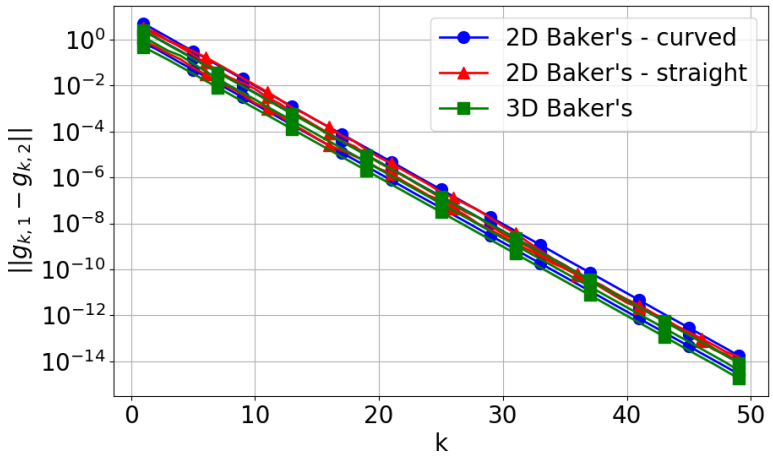

Figure 3-16: Relation between $\|g_{k,1} - g_{k,2}\|$ and the time step k in the semilogarithmic scaling. This plot contains nine curves of three different colors. Each color corresponds to a different map: 2D Baker's map with curved unstable manifolds (blue), 2D Baker's map with straight unstable manifolds (red), and 3D Baker's map with $s_1 = 0$, $s_2 = 0.9$, $s_3 = 0.1$ (green). In case of the 2D Baker's map, the parameter values are the same as those in Figure 3-14.

three different chaotic models is plotted in Figure 3-16. The g function is generated using Algorithm 1 for the 2D Baker's map with $m = 1$ (Eq. 3.69), as well as the 3D Baker's map with $m = 2$ (Eq. 3.84).

We observe the norm-versus-k relation is linear in the semilogarithmic scale, which clearly indicates an exponential convergence of our algorithm if applied to the Baker's map. This result implies that a relatively small number of steps ($k \approx 50$) is required to obtain the machine-precision value of the norm. Note also that the choice of a trajectory (x_0) or model has a negligible effect on the error.

3.8.3 Numerical example: Monte Carlo integration

To validate Algorithm 1, we consider a square-integrable function $f(x) \in L^2(\mu)$ and integrate it with respect to the SRB measure μ using a Monte Carlo procedure. By the Central Limit Theorem, this integral can be approximated by taking the average of the sample distributed according to μ, while the approximation error is upper-

bounded by $\sqrt{\mathrm{Var}(f)/N}$, i.e.,

$$\left| \int_M f(x)\, d\mu(x) - \frac{1}{N} \sum_{k=0}^{N-1} f(x_k) \right| \leq C \sqrt{\frac{\mathrm{Var}(f)}{N}}, \tag{3.85}$$

where $C > 0$ and $x_{k+1} = \varphi(x_k) \in M$. Therefore, by generating a sufficiently long trajectory and evaluating f at every point along it, we gradually approach the sought-after solution. Motivated by particular applications of the SRB density gradient function (see Section 3.6), we consider $f(x) := \sum_{j=1}^{m} \partial_{Q^{(:j)}} v(x)$, where $v(x) : M \to \mathbb{R}$ is some smooth function. In other words, we strive to integrate a sum of m directional derivatives along m-dimensional unstable manifolds of the scalar function $v(x)$. Note integrals of this type are critical in the sensitivity computation using, for example, the general S3 method [32, 31]. Thanks to the partial integration (see Eq. 3.50-3.56), we can apply Monte Carlo to two alternative versions of the same integral, since

$$\int_M \sum_{j=1}^{m} \partial_{Q^{(:j)}} v(x)\, d\mu(x) = I = -\int_M \sum_{j=0}^{m} g^{(j)}(x)\, v(x)\, d\mu(x). \tag{3.86}$$

Using this equation, we validate Algorithm 1 for g by comparing numerical approximations of the LHS and RHS. Due to its trajectory-driven structure, Algorithm 1 is naturally compatible with the Monte Carlo procedure.

Two different maps shall be tested. First, we shall consider the 2D Baker's map (Eq. 3.69) with $s_4 = 0.4$ and $s_1 = s_2 = s_3 = 0$. As illustrated in Figure 3-14, its unstable manifolds are curved and therefore the simplified version of the recursion for g (Eq. 3.68) cannot be used. In this particular case, q has in fact two nonzero components. Indeed, we numerically estimate that

$$\max_{k \in 1,2,\ldots,N} \arctan \left| \frac{q_k^{(2)}}{q_k^{(1)}} \right| \approx 0.24\,\mathrm{rad} \approx 14^\circ,$$

which is consistent with the illustration of unstable manifolds in Figure 3-14. The second map is the 3D Baker's map (Eq. 3.84) with $s_1 = 0, s_2 = 0.9$, $s_3 = 0.1$. One can easily verify unstable manifolds of this map are flat surfaces aligned with

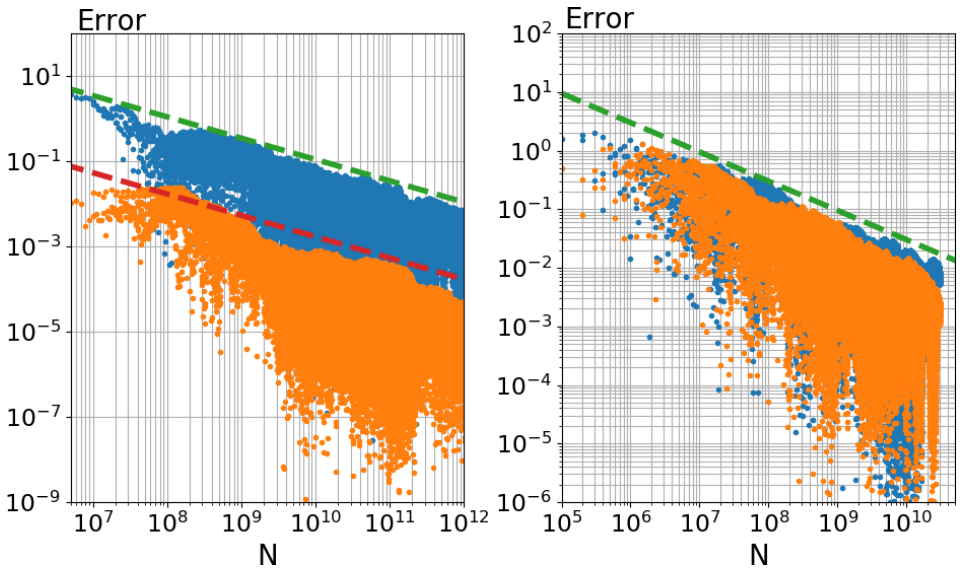

Figure 3-17: Error of the Monte Carlo approximation of the LHS (blue dots) and RHS (orange dots) of Eq. 3.86 versus the amount of data N used. Left: 2D Baker's map. Here, we compute the relative error with respect to the reference value -1.05335809 (which equals the approximation of the RHS integral at $N = 10^{13}$) for $v(x) = \sin(x^{(1)}) \exp(x^{(2)})$. Right: 3D Baker's map. Here, we compute the absolute error with respect to the reference value of 0 for $v(x) = \sin(x^{(1)}) \sin(3/2\, x^{(2)})\, x^{(3)}$. The dashed lines represent the slope $-1/2$ in the logarithmic scaling.

the $x^{(1)}$–$x^{(2)}$ plane. These expanding surfaces could be curved by adding an $x^{(3)}$-dependent perturbation term to the third component of the map.

Figure 3-17 includes results of the integration test. Our primary conclusion is that the Monte Carlo approximations of the LHS and RHS of Eq. 3.86 approach each other as $N \to \infty$ with the rate $\mathcal{O}(1/\sqrt{N})$, which directly confirms the correctness of Algorithm 1. Recall we require g to regularize the linear response formula, as it involves derivatives of strongly-oscillatory functions (see Section 3.6). The examples presented in this section, however, include mildly-oscillatory functions $v(x)$ with derivatives that behave similarly (note they involve a combination of trigonometric, exponential and linear functions). Nevertheless, we observe significantly smaller errors of the RHS approximation in the 2D Baker's map case. Note the approximation error of Monte Carlo integration also depends on the variance of the integrand, which can be upperbounded by a quantity proportional to the $L^2(\mu)$-norm of the SRB density gradient g, denoted by $\|g\|_{L^2(\mu)}$. Indeed, $\|g\|_{L^2(\mu)}$ equals $\mathcal{O}(10^{-2})$ and $\mathcal{O}(10^{1})$

for the 2D and 3D Baker's map, respectively. This explains the significantly better performance of the Monte Carlo procedure in the former case. Therefore, if $\|g\|_{L^2(\mu)}$ exists and is sufficiently small, Monte Carlo integration might be significantly cheaper if applied to the regularized integrals of this type, regardless of the behavior of $v(x)$. If g is not even Lebesgue-integrable, i.e. $g \notin L^1(\mu)$, the integrals in Eq. 3.86 do not converge, as showed in [159].

3.9 Summary

Ruelle's linear response formula is fundamental in the construction of numerical methods for sensitivity analysis of n-dimensional hyperbolic chaotic systems. Its original form, however, is impractical for direct computation due to the presence of derivatives of composite functions that grow exponentially in time. Fortunately, it is possible to easily regularize this expression through partial integration. In case of nonuniform measures describing the statistics of chaos, the by-product of the integration by parts, per the generalized fundamental theorem of calculus, involves the SRB density gradient g defined as the directional derivative of conditional SRB density on m-dimensional unstable manifolds. Computation of g is the price that must be paid for a computable version of Ruelle's formula.

Using the measure-based coordinate parameterization, the time evolution of the measure gradient is rigorously derived by applying the measure preservation property, differentiating the coordinate charts with the chain rule on smooth manifolds. Indeed, g can be computed in a recursive manner by solving a set of $\mathcal{O}(m)$ first- and $\mathcal{O}(m^2)$ second-order tangent equations, as well as step-by-step QR-factorization and inversion of $n \times m$ and $m \times m$ matrices, respectively. While the total cost of approximating g at N consecutive points along a trajectory is $\mathcal{O}(Nn^3m^2)$ in the worst-case scenario, the actual computational cost may scale linearly with the dimension of the system in many real-world models due to their sparse structure. Moreover, this procedure requires storing $\mathcal{O}(m^2)$ n-dimensional vectors only from the current time step to advance the iteration in time. Therefore, in terms of the hardware requirements, our algorithm

would definitely be a reasonable choice for high-dimensional physical systems since $m \ll n$.

The algorithm we propose is compatible with existing methods for sensitivity analysis that stem from the linear response theory, including the space-split sensitivity (S3) and FDT-based methods. Many of them approximate sensitivities through an ergodic-averaging Monte Carlo procedure and require knowledge of the directional derivative of conditional SRB measures. Moreover, g can be used to assess the differentiability of statistical quantities in hyperbolic systems, which is a recurring theme in theoretical studies of chaos. Thus, we believe our method provides a new major tool for both rigorous analysis and applied studies of large chaotic systems.

3.10 Appendices

This section includes supplementary material supporting certain arguments presented in this chapter.

3.10.1 Applying the simplified recursive formula for SRB density gradient to 1D non-injective maps

Throughout this chapter, we assume φ is an invertible map. Based on this assumption, we directly use the measure preservation property to derive a recursive formula for g, including the simplified version for maps with straight unstable manifolds, as described in Section 3.7.2. However, in the literature, one can find several one-dimensional maps such as the sawtooth/Bernoulli map [158], cusp map [121], logistic map [189], onion map [159], tent map [11], and so forth. All of them are scientifically relevant, as they represent some simplified physics or feature interesting mathematical properties. However, most of them are non-injective, which violates the basic assumption of our derivation. In this section, however, we argue that Eq. 3.68 can still be used to compute g for such maps.

Many of the popular 1D chaotic maps (such as those listed above) are two-to-one.

Thus, we assume φ satisfies this condition; however, the argument we present can be naturally extended to other types of surjection. Let us also assume, without loss of generality, $\varphi : [0,1] \to [0,1]$ and φ is monotonic in $[0, 0.5)$ and $(0.5, 1]$. Let us now define a two-dimensional analog of φ, denoted by $\varphi_{2D} : [0,1]^2 \to [0,1]^2$ and satisfying

$$x_{k+1} = \varphi_{2D}(x_k) = \begin{bmatrix} \varphi(x_k^{(1)}) \\ x_k^{(2)}/2 + 0.5\lfloor 2x_k^{(1)} \rfloor \end{bmatrix}. \tag{3.87}$$

Note φ_{2D} is invertible and resembles the 2D Baker's map (see Eq 3.69). The invertibility is guaranteed by adding the floor function in $\varphi_{2D}^{(2)}$. Analogously to the 2D/3D Baker's map, here the discontinuity point is located at $x^{(1)} = 0.5$, which means that the value of 0.5 is added to $x_k^{(2)}/2$ if $x_k^{(1)} > 0.5$. If the monotonicity breaking point was different, then the coefficients of the floor function would need to be modified accordingly. One of the main messages of this example is to point out that any surjective 1D map can be represented as a higher-dimensional invertible map with one positive Lyapunov exponent.

Note 1D unstable manifolds of φ_{2D} are aligned with the first phase space coordinate, per the argument given in Section 3.7.2. Thus, its SRB distribution is similar to the one of Baker's map presented in Figure 3-14. Note also that the horizontal deformation of the trajectory of φ_{2D} is solely determined by φ. This implies that the SRB distribution of φ is in fact an integral of SRB distributions of φ_{2D} restricted to single unstable manifolds over all values of $x^{(2)}$. In other words, φ_{2D} scatters the SRB measure of φ (which is supported on [0,1]) over an infinite set of vertically stacked intervals $[0, 1]$ (which geometrically coincide with unstable manifolds of φ_{2D}). This further implies the SRB density of φ equals the SRB distribution of φ_{2D} integrated with respect to the vertical (second) coordinate.

In case of the map defined by Eq. 3.87, the simplified recursive formula for g can be expressed in terms of phase space derivatives of φ (see Section 3.7.2 for the derivation),

$$g(\varphi(x)) = \frac{g(x)}{\varphi'(x)} - \frac{\varphi''(x)}{\varphi'(x)^2}. \tag{3.88}$$

Note that Eq. 3.88 is identical as Eq. 2.36. Here, the prime symbol ($'$) denotes differentiation with respect to phase space. Let $\rho(x)$ be the SRB density of φ. The $g(x)$ function that satisfies Eq. 3.88 is not the SRB density gradient of φ, defined as $g^{\varphi}(x) := \rho'(x)/\rho(x)$. According to our discussion in Section 3.7.2, $g(x)$ is in fact a conditional SRB density gradient of φ_{2D} associated with the unstable manifold parameterized by $x^{(2)}$. However, as we discussed in the previous paragraph, the SRB measure of φ can be computed by integrating "slices" of the SRB measure of φ_{2D} parallel to $x^{(1)}$. This implies that, given the definition of the SRB density gradient, $g^{\varphi}(x)$ can be computed by Lebesgue-integrating the SRB density gradients obtained in the above iteration along the vertical axis.

In practice, to construct a trajectory-based algorithm for g^{φ}, we can directly use the recursion in Eq. 3.88. The algorithm we propose is the following. Divide the phase space $[0, 1]$ in $K \in \mathbb{Z}^+$ bins of equal width. Generate a sufficiently long sequence $\{g_0, g_1, g_2, ...\}$ using Eq. 3.88 starting from a random initial condition g_0. For each bin, take the average of the members of the sequence that correspond to one bin. Based on our discussion above, the obtained average value converges to g^{φ}. This algorithm in fact provides a piecewise constant approximation of g^{φ}.

To verify our argument, we present a numerical experiment in which we apply the algorithm to two different 1D maps, the sawtooth map and onion map. Both of them are two-to-one and piecewise smooth. Figure 3-18 shows raw values of the sequence $\{g_0, g_1, g_2, ...\}$ obtained using Eq. 3.88, their averaged values, and finite-difference (FD) approximation of g^{φ} using empirical SRB densities of these maps. We observe there is a good agreement between the averaged values and FD approximations in both cases.

Finally, we perform the relative error convergence test of the averaged values with respect to the trajectory length N. We focus on two different bins and compute the relative error with respect to a reference value generated using significantly more samples. Our results generated for the sawtooth map are shown in Figure 3-19. As expected, the error decays and is upperbounded by $\mathcal{O}(1/\sqrt{N})$, which is a consequence of the Lebesgue-integration (or, equivalently, weighted averaging) of (conditional)

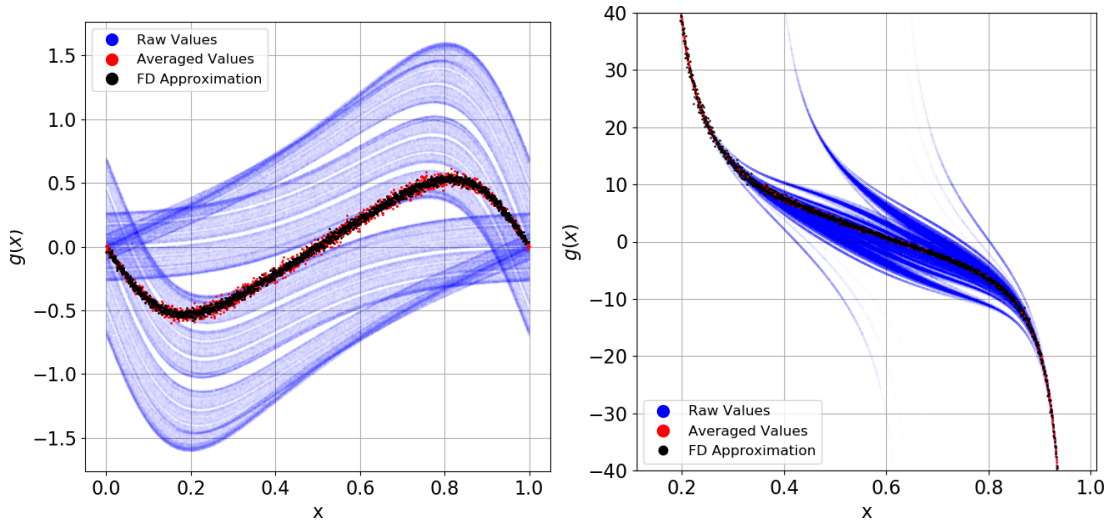

Figure 3-18: SRB density gradient generated for the sawtooth map $x_{k+1} = 2x_k + s\,\sin(2\pi\,x_k) \bmod 1$ at $s = 0.1$ (left) and the onion map $x_{k+1} = 0.97\sqrt{1-|1-2x_k|^\gamma}$ at $\gamma = 0.4$ (right). The averaged values (red dots) were computed by averaging the raw values (blue dots) in each of 2048 bins. The FD Approximation data points represent the central finite difference approximation of the SRB density gradient using the definition of g and empirically computed SRB densities. We generated a trajectory of length $N = 10^6$ to compute the raw/averaged values of g.

SRB density gradients. This example shows that a trajectory of minimum length $N = 10^9$ should be generated in order to obtain an approximation with a relative error smaller than 1%.

3.10.2 Probing the hyperbolicity of the Baker's map

Hyperbolicity guarantees the tangent space can be decomposed into two $D\varphi$-invariant subspaces, where one is asymptotically expanding (unstable), while the other one is asymptotically contracting (stable). If the expansion/contraction is uniform, then such systems are uniformly hyperbolic. Hyperbolicity is the major assumption for the dynamical systems we consider in this chapter. Indeed, if the system is hyperbolic and has absolutely continuous conditional measures on unstable manifolds, then the SRB measure exists [40]. It is not always possible to analytically verify that a particular map is hyperbolic. Fortunately, there exist numerical procedures allowing for an efficient assessment of hyperbolicity [90]. Most of them test the two basic criteria of hyperbolicity: 1) No zero LEs, and 2) Strict separation of the stable and unstable

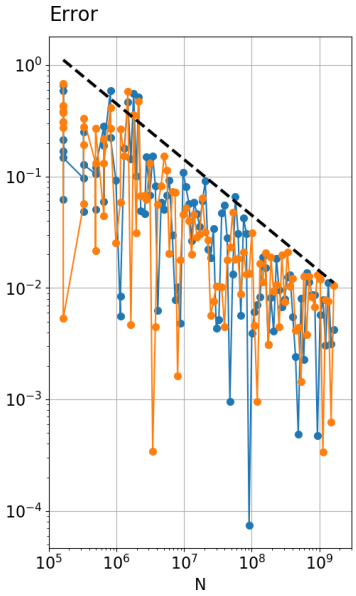

Figure 3-19: Relative error of the approximation of $g^{\varphi}(x)$ versus the trajectory length N. The error was computed for the sawtooth map at $s = 0.1$ at two phase space coordinates, $x \approx 0.4$ (blue curve) and $x \approx 0.6$ (orange curve). All error values were computed with respect to the reference value generated using $N = 3.3 \cdot 10^{11}$ samples. The reference dashed line represents the slope $-1/2$ in the logarithmic scaling.

subspaces. Here, we apply the method proposed in [89], which computes the basis vectors of the two subspaces and approximates the smallest angle between them at different points of the manifold. If any of these angles is close to zero, then the stable and unstable subspaces are (almost) tangent, which implies the systems is likely to be non-hyperbolic. In Figure 3-20, we compute the PDF of $d \in [0, 1]$, which is a normalized quantity associated with the smallest principal angle between the stable and unstable subspace (our d equals $k!\, d_k$; see the above reference for a rigorous definition of d_k). If the distribution is evidently separated from the origin ($d = 0$), then it is highly likely there are no tangencies between the two subspaces. We observe the normalized parameter d is highly unlikely to drop below the value of 0.97. As a by-product of the applied algorithm, we computed the spectrum of Lyapunov exponents (alternatively, one can use Benettin et al.'s algorithm [15]). The LEs approximately equal: $0.69 \approx \log(2)$, -0.69 (2D Baker's with straight unstable subspaces), $0.69 \approx \log(2)$, -0.71 (2D Baker's with curved unstable subspaces), $1.09 \approx \log(3)$, $0.69 \approx \log(2)$, -1.16 (3D Baker's). Although a small change in the parameter

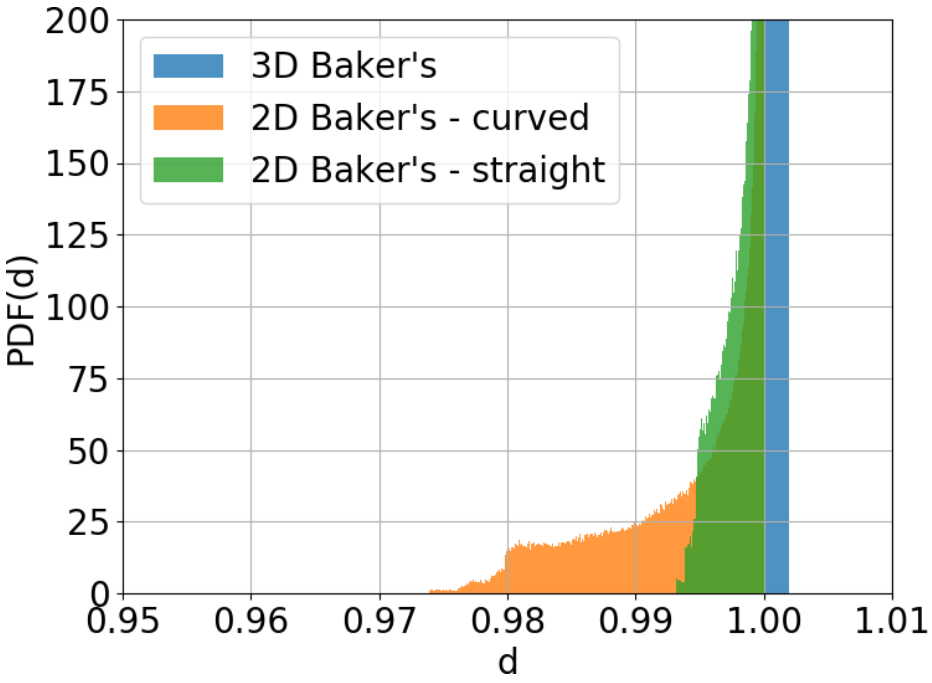

Figure 3-20: Numerical hyperbolicity tests of the 2D and 3D Bakers map. The parameter values are the same as the ones used is the numerical examples in Section 3.7 (2D Baker's) and Section 3.8 (3D Baker's). To generate the PDF, we computed $N = 10^6$ samples of d along a trajectory.

value does not significantly impact the LE values, it may move the PDF of d closer to the origin. Based on the empirical evidence presented in this section, we conclude the 2D/3D Baker's map is clearly hyperbolic at the chosen parameter values.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 4

Assessment of smooth and rough parameter dependence of statistics in chaotic dynamical systems

Modern linear response algorithms generally assume that statistics or long-time averages implied by a chaotic system smoothly depend on input parameters. An accurate approximation of the underlying linear response is critical is several fields of science and engineering. The purpose of this chapter is to develop a generalizable numerical method for the assessment of the differentiability of statistics with respect to controllable parameters in chaotic systems. We numerically show that the existence of the linear response strictly depends on the Lebesgue-integrability of the SRB measure gradient. We demonstrate the performance of the proposed method in determining the differentiability of statistics using low-dimensional systems. The chosen examples feature both smooth and rough statistical behavior with respect to the input.

The content of this chapter was published in *Communications in Nonlinear Science and Numerical Simulation* by *Elsevier* [159].

4.1 Does the linear response really exist?

Ruelle [148, 150] showed that the statistics-parameter relation is guaranteed to be smooth in uniformly hyperbolic systems. Gallavotti and Cohen [59] conjecture that several chaotic systems encountered in science and engineering behave as if they were uniformly hyperbolic. It means that certain systems with locally non-uniformly hyperbolic behavior may still feature smooth ergodic measures and/or exponentially fast mixing [103]. This conjecture, known as the chaotic hypothesis, is also supported by empirical evidence involving well-known physical models, including ODE models, such as Lorenz 96 [163], as well as PDE models for turbulence governed by Kuramoto-Sivashinsky equation [21, 163] and the 3D Navier-Stokes equation [126, 20]. The linear response theory was successfully applied to various climate prediction models used to better understand fluctuations of climate properties [114, 3, 105]. However, some models used in climate modeling and geophysical fluid dynamics indicate that the chaotic hypothesis cannot always be applied. For example, a rough statistics-parameter relationship has been observed in an El-Niño Southern Oscillation climate model in [35]; other work describes a similar phenomenon in a barotropic model of atmosphere of the northern hemisphere with forcing and dissipation [70]. In [189], the notion of the existence of linear response is addressed from the perspective of statistical mechanics, which was motivated by earlier work related to stochastic systems [73]. The authors present a simple multidimensional model whose macroscopic variables vary differentiably with respect to the input, as opposed to its individual microscopic subsystems. This clearly indicates a high complexity of the posed question when considering systems with several degrees of freedom and multi-scale dynamical interactions. Recently published numerical evidence [187] presents a high-dimensional system with a clearly non-smooth statistical behavior of its large-scale dynamical structures. which The authors of [188] warn that a direct numerical reproduction of the statistics-parameter curve is sensitive to perturbations in data. Consequently, a more detailed information about system's ergodic measure should also be computed. Chekroun *et al.*'s work [35] establishes a relation between the spectral gap of the

transfer operator [10] and the smoothness of statistical quantities in the parametric space. This dependence is used to construct a numerical framework for the assessment of the differentiability of statistics. However, the authors of [170] argue that indicators of attractor crisis, expressed in terms of a correlation function, may not be generalizable to high-dimensional systems. We also acknowledge that a generalization of the Fluctuation-Dissipation Theorem (FDT) was utilized in [69] to study the response attributes of an atmospheric general circulation model, and to verify the applicability of the linear response theory for that particular model; similar results obtained with FDT-based problem-dependent methods can also be found in [39, 95].

This chapter offers an alternative computational framework assessing the existence of linear response. We present a mathematical argument based on Ruelle's formula indicating a relation between the distribution of the SRB measure gradient and smoothness of long-time averages. Our numerical experiments further demonstrate that when the SRB measure gradient is Lebesgue-integrable, the linear response holds and vice versa. To validate our results, we develop an auxiliary numerical method approximating the Hölder exponent of a one-dimensional function. Our numerical framework has two attractive features: i) solely based on one-step trajectory-driven tangents equations, and ii) easily generalizable to multi-dimensional systems with an arbitrary number of positive Lyapunov exponents.

4.2 Onion map: Example of simple chaos with statistically rough behavior

At the beginning of our discussion, we present a simple chaotic map that features both smooth and non-smooth (rough) behavior. In that map, the degree of smoothness of statistical quantities strictly depends on the value of its parameters. Let us consider a one-dimensional map, $\varphi : [0, 1] \to [0, h]$, defined as follows,

$$x_{k+1} = \varphi(x_k; \gamma, h) = h\sqrt{1 - |1 - 2x_k|^{\gamma}}. \tag{4.1}$$

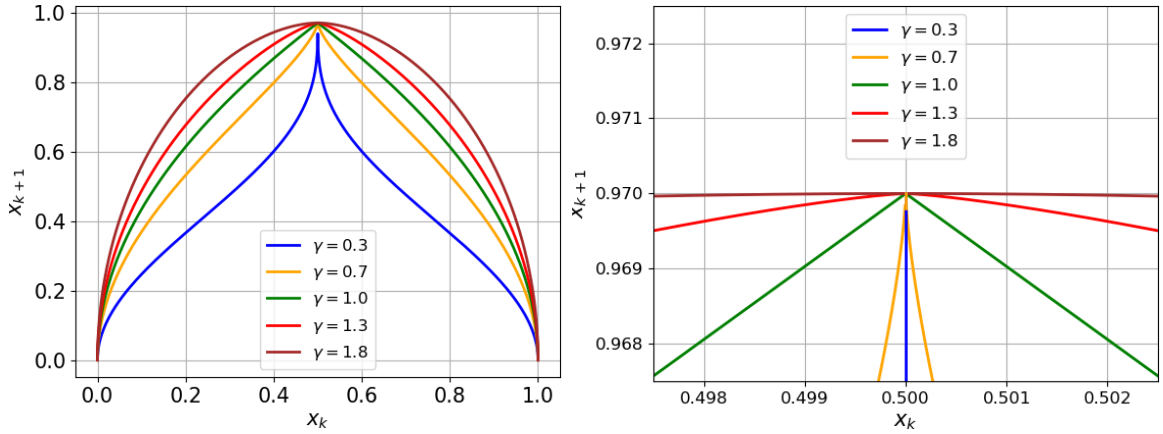

Figure 4-1: Illustration of the onion map at $h = 0.97$ and its dependence on γ. The right-hand side plot zooms in the region in the vicinity of the tip.

where $\gamma > 0$, $1 \geq h > 0$ are map parameters. To simplify the notation, we will skip the parameters in the argument list, i.e., $\varphi(x_k) = \varphi(x_k; \gamma, h)$. Figure 4-1 depicts Eq. 4.1 at some selected parameter values. Due to its characteristic shape, this map will be further referred to as the onion map. While the proportionality parameter h only affects the range of φ, the exponent γ has a significant impact on the function shape in the vicinity of the tip located at $(0.5, h)$. If $\gamma < 1$, the tip is sharp, i.e., the derivative $\varphi'(0.5)$ does not exist, and its shape resembles the cusp map, as defined and illustrated in [19]. The cusp map has been used in modeling as a one-dimensional simplification of the Lorenz 63 system [121]. We observe the tip blunts when γ gets larger than 1, and the shape of φ converges to the well-known logistic map [120] as γ approaches the value of 2. Given the cusp map and logistic map feature smooth and non-smooth statistical behaviors [19], the onion map, which combines both of them, is a perfect example of a map with varying regularity of statistical quantities. In our numerical examples, we fix the value of h to 0.97 and consider different values of γ. One can easily verify that the only Lyapunov exponent (LE) λ, defined by Eq. 2.4, is always positive for the onion map if $h = 0.97$ and $\gamma \in [0.15, 1.85]$. If we change the value of h, then the range of γ for which $\lambda > 0$ is slightly different. A positive value of the Lyapunov exponent implies *chaotic* behavior of the map reflected by the *butterfly effect*, i.e. strong sensitivity to the initial conditions. LE measures rate of separation of two trajectories of a chaotic map and its value depends only on the parameter.

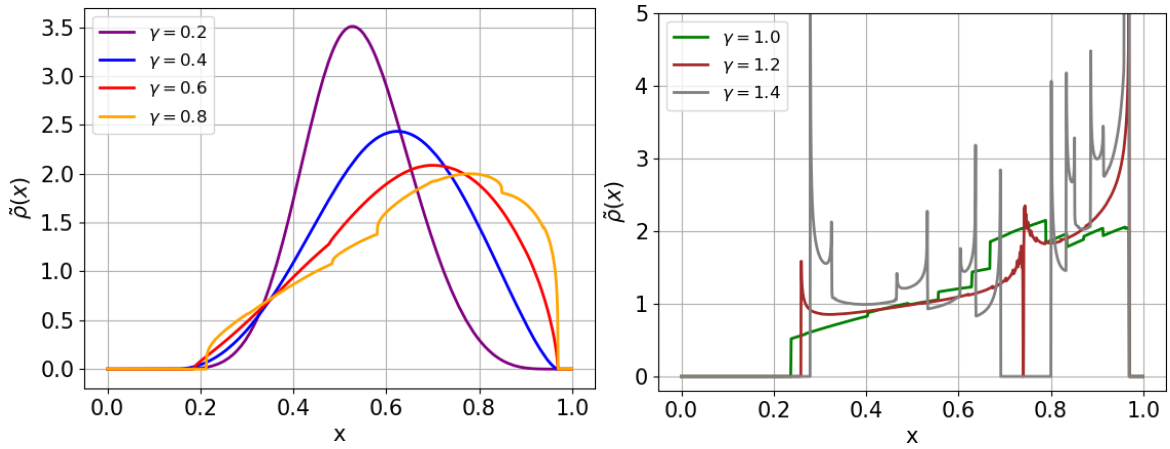

Figure 4-2: Empirical density distribution $\tilde{\rho}(x)$ generated for the onion map (Eq. 4.1) at $h = 0.97$. To generate $\tilde{\rho}(x)$, we divided the domain $x \in [0, 1]$ into $K = 2048$ bins of equal width, counted the number of times the trajectory passes through each bin. We used $N = 41,943,040,000$ samples to per histogram. The obtained histogram is normalized, through the multiplication by K/N, to satisfy the axiom of unit area.

As thoroughly described in Chapter 3, a function critical in the analysis of chaotic systems is the Sinai-Ruelle-Bowen density ρ, which contains statistical information of the dynamics described by φ [194, 195, 41, 158]. Intuitively, ρ can be viewed as the likelihood of the trajectory passing through a non-zero-volume region of the manifold and, if normalized, ρ can be viewed as a probability density function. In case of one-dimensional maps defined on $[0, 1]$, the SRB density ρ is a function that maps $[0, 1]$ to the set of non-negative real numbers, which satisfies the unity axiom, i.e., $\int_0^1 \rho(x)\, dx = 1$. Figure 4-2 shows empirical densities $\tilde{\rho}$ generated for the onion map at different values of the exponent γ. We observe the empirical density distribution is smooth if $\gamma \leq 0.4$. When the exponent γ becomes higher, but is still no larger than 1, the approximated density is clearly bounded, but have some non-smooth regions. If $\gamma \geq 1$, the $\tilde{\rho}$ distribution features discontinuous regions. The onion map is a purely expanding map if $0.24 < \gamma < 0.99$, and thus SRB density exists in that parameter regime [100]. Therefore, the observed empirical densities on the left-hand side plot of Figure 4-2, marked with blue, red and yellow, are rigorously guaranteed to be approximations of ρ.

Although no rigorous proof is available, we assume the onion map is ergodic, which can be justified by solid numerical evidence for onion map-like systems [19].

Ergodicity in fact means the statistics of the onion map do not depend on initial conditions. Moreover, this property implies that the Ergodic Theorem holds, and thus we can directly use the stationary density to compute the long-time averages of the onion map. In particular, the theorem ensures that an infinite time average of some quantity of interest J is equal to the expected value of the same quantity computed with respect to the density distribution. Mathematically, it means that

$$\langle J \rangle = \lim_{N \to \infty} \frac{1}{N} \sum_{k=0}^{N-1} J(x_k) = \int_0^1 J(x) \rho(x) dx \tag{4.2}$$

always holds. It is assumed that J is an integrable bounded function and it does not depend on the map parameter. The origins of the first assumption will be explained in Section 4.3. Furthermore, the dependence of the objective function on the parameter is not considered in this chapter, as it does not impose extra mathematical complexity in the computation of sensitivities and is irrelevant in the context of our analysis. Therefore, two critical properties can be further inferred from Eq. 4.2. First, the map statistics $\langle J \rangle$ solely depends on its parameter γ and, second, the smoothness of statistics strictly depends on the smoothness of the density distribution. In our numerical test, we set $J(x) = \delta_c^\epsilon(x)$, $c \in [0, h]$, where $\delta_c^\epsilon(x)$ is an indicator function, i.e. $J(x) = 1$ for all $x \in x_c^\epsilon := [c - \epsilon/2, c + \epsilon/2]$, and $J(x) = 0$ otherwise. Note with this particular choice of the quantity of interest, the long-time average equals the density distribution itself evaluated at c in the limit $\epsilon \to 0$, i.e. $\langle J \rangle = \rho(c)$ if ϵ is infinitesimally small. For any ϵ such that $x_c^\epsilon \subset [0, 1]$, $\langle J \rangle$ equals the integral of ρ over x_c^ϵ. Note also that for any Riemann-integrable $J(x)$, the statistics can be easily computed using a numerical integration scheme by virtue of Eq. 4.2 if ρ is available. We emphasize that the indicator function can be viewed as an approximation of an infinitely-smooth steep function. In the context of numerical computation of Eq. 4.2, it does not matter whether we choose a smooth or non-smooth version of the indicator function. We pick the latter for convenience, as it exempts us from applying numerical integration at every single value of γ. Figure 4-3 illustrates the relationship between $\langle J \rangle$ and γ, for two different values of c.

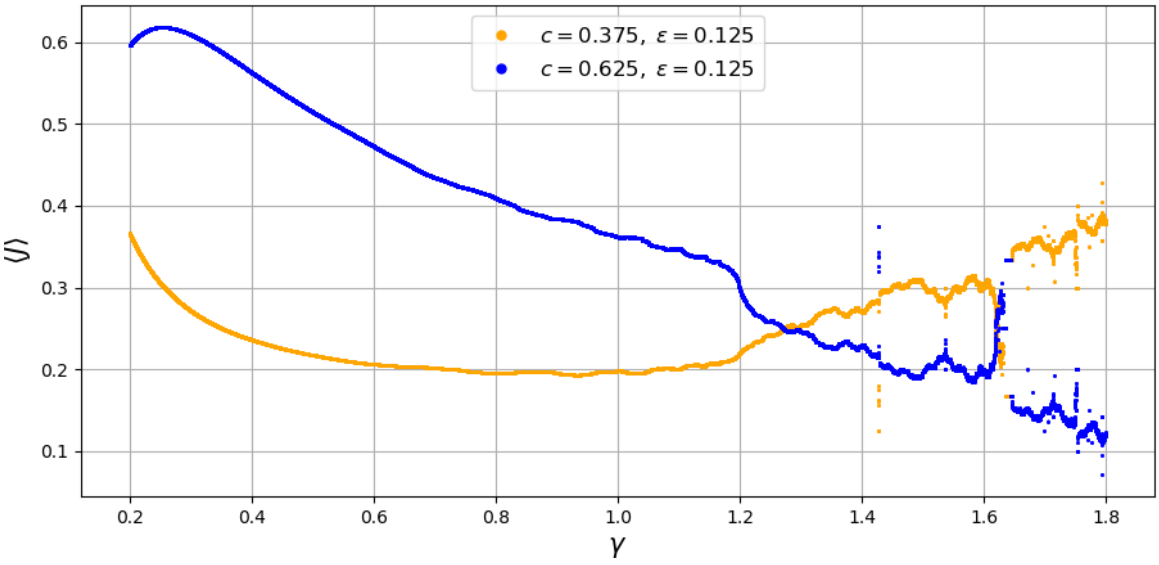

Figure 4-3: Relationship between the long-term average and the exponent γ for the onion map (Eq. 4.1) at $h = 0.97$ with $J(x) = \delta_c^\epsilon(x)$. To generate this plot, we computed density distributions at a uniform grid of 16,001 different values of γ between 0.2 and 1.8. For each value of γ, we run 10 independent simulations using $N = 41,943,040,000$ samples per simulation. In the calculation of the density, we divided the domain $x \in [0, 1]$ into $K = 4$ bins of equal width (see Figure 4-2 for reference).

In terms of the function smoothness, we observe a similar trend in both Figure 4-2 and Figure 4-3. In particular, if γ increases, both the density function and long-time average become more oscillatory and even discontinuous. This result is consistent with the study in [19], where the relationship between the smoothness of the statistics and smoothness of the density distribution has been justified analytically using the Frobenius-Perron operator, which belongs to the class of Markov operators and describes the evolution of the SRB density [45]. The authors of [19] notice that if the Frobenius-Perron operator is well-conditioned and the density function is differentiable in phase space (with respect to x in 1D), then $\partial\rho/\partial\gamma$ must be bounded. This observation is critical for the existence of $d\langle J\rangle/d\gamma = \int_0^1 J(x)\frac{\partial\rho}{\partial\gamma}(x)dx$, which is the sought-after quantity in sensitivity analysis.

Motivated by the above discussion, we strive to find a computable and generalizable mathematical criterion for the existence of $d\langle J\rangle/d\gamma$. In particular, our purpose is to identify a condition that can be translated to an efficient numerical method, and

is applicable to higher-dimensional chaotic systems.

4.3 SRB measure gradient as an indicator of the differentiability of statistics of 1D chaos

The main focus of this section is to highlight the significance of the SRB density (measure) gradient function, denoted by g, in the context of the differentiability of statistics in one-dimensional chaotic maps. This function is a fundamental ingredient of the derivative of statistics. Here, we consider 1D chaotic maps in which case g is defined as follows,

$$g(x) = \frac{d \log \rho}{dx}(x) = \frac{\rho'(x)}{\rho(x)}. \tag{4.3}$$

That is, the density gradient g is the relative rate of change of the SRB density at each point on the 1D manifold (see Chapter 2). Throughout this section, we use the prime symbol ($'$) to indicate differentiation with respect to phase space. We assume $\varphi : [0,1] \to [0,1]$ is a 1D, invertible, ergodic, C^3 map with a positive LE. Let J be a smooth observable whose expectation with respect to the SRB density or equivalently, the infinite-time average starting from almost everywhere, is denoted $\langle J \rangle$. In this case, Ruelle's *linear response formula* [148, 150], which is a closed-form expression for the parametric derivative of $\langle J \rangle$, is given by

$$\frac{d\langle J \rangle}{d\gamma} = \frac{d}{d\gamma} \int_0^1 J(x)\ \rho(x)\, dx = \sum_{k=0}^{\infty} \int_0^1 f(x) \left(J \circ \varphi_k\right)'(x)\ \rho(x)\, dx, \tag{4.4}$$

where $f := \partial\varphi/\partial\gamma \circ \varphi^{-1}$ is the parameter perturbation. The subscript notation is used to denote the number of times a map φ is applied i.e., $\varphi_0(x) = x$ and $\varphi_k(x) = \varphi(\varphi_{k-1}(x))$ for any state vector x, while the inverse of the map is indicated using the conventional notation, i.e., φ^{-1}. Integrating the RHS of Eq. 4.4 by parts leads to an alternative expression for the sensitivity,

$$\frac{d}{d\gamma} \int_0^1 J(x)\ \rho(x)\, dx = -\sum_{k=0}^{\infty} \int_0^1 \Big(g(x)\ f(x) + f'(x) \Big) (J \circ \varphi_k)(x)\ \rho(x)\, dx, \tag{4.5}$$

which provides a direct relation between the derivative of the long-time average and the density gradient function g (see Chapter 2 for the derivation of Eq. 4.5). Eq. 4.5 is in fact a one-dimensional version of the space-split sensitivity (S3) formula, originally derived and computed in [32]. A version of S3, which shows the relation $d\langle J\rangle/d\gamma$ vs. g for multi-dimensional systems with one positive LE, is discussed in Section 4.5.1.

The k-time correlation between two observables ϕ and ψ is given by

$$\begin{aligned} C_{\phi,\psi}(k) = &\int_0^1 \phi \circ \varphi_k(x)\, \psi(x)\, \rho(x)\, dx - \\ &\Big(\int_0^1 \phi(x)\, \rho(x)\, dx\Big)\Big(\int_0^1 \psi(x)\, \rho(x)\, dx\Big). \end{aligned} \tag{4.6}$$

Using this definition and Eq. 4.5, we conclude that

$$\frac{d}{d\gamma}\int_0^1 J(x)\, \rho(x)\, dx = -\sum_{k=0}^{\infty} C_{J,\, gf+f'}(k). \tag{4.7}$$

Uniform hyperbolicity guarantees that $C_{\phi,\psi}(k)$ decays exponentially with $k > 0$ for any pair of observables (ϕ, ψ) at a uniform rate [148, 36] within the respective function class. Thus, there exists a positive constant c such that

$$C_{\phi,\psi}(k) \sim \mathcal{O}(e^{-ck}).$$

Eq. 4.7 is valid due to periodic boundary conditions. Indeed, notice that the k-term of Eq. 4.5 is in fact a k-time correlation between J and $h := gf + f'$ because

$$\int_0^1 \Big(g(x)\, f(x) + f'(x)\Big)\, \rho(x)\, dx = \int_0^1 \Big(\rho(x)\, f(x)\Big)'\, dx = 0. \tag{4.8}$$

Similar divergence terms arise from partial integration in the case of multi-dimensional systems with more than LE[1] also vanishes (see Section 4.5.1). One can show that the boundary terms always vanish. Therefore, we conclude that if $C_{J,h}(k)$ is summable,

[1]This procedure can be performed only on the unstable manifold because, in general, the SRB density is absolutely continuous along that manifold only.

Ruelle's formula converges and the linear response is bounded. Now, the question is about the summability of $C_{J,h}(k)$. Our goal is to use the above conclusion to identify a computable criterion for the differentiability of statistics. Let us consider again the the k-th term of the linear response,

$$C_{J,h}(k) = C_{J,f'}(k) + C_{J,\,fg}(k). \tag{4.9}$$

In our analysis, we assume $J \in C^1$ and $f \in C^2$. Thus, the sequence of the time correlations $C_{J,f'}(k)$ decays exponentially fast. Typically, in the context of real-world simulations, J represents a physical quantity, such as force or temperature, which are smooth functions. The function f and its derivative f' are also smooth in several physical ODE/PDE models [110, 166, 21, 163].

Notice that if $C_{J,f'}(k)$ is absolutely summable, then the existence of the linear response would depend only on the absolute summability of the remaining time correlation, $C_{J,\,gf}(k)$. Consequently, there might exist weaker conditions for g that guarantee the convergence of Ruelle's series, regardless on the choice of the objective function and input perturbation. Without loss of generality, let us assume $\langle J \rangle = 0$ to eliminate the second term of 4.6. This assumption does not violate the consistency of our analysis because $d/d\gamma \, \langle J \rangle = d/d\gamma \, \langle J - \langle J \rangle \rangle$. Thus, there exists a constant $c > 0$ such that for all $K \in \mathbb{Z}^+$

$$\begin{aligned} &\left| \sum_{k \le K} \left(\int_0^1 J \circ \varphi_k(x) \; g(x) \; f(x) \; \rho(x) \; dx + \int_0^1 J \circ \varphi_k(x) \; f'(x) \; \rho(x) \; dx \right) \right| \\ &\le \sum_{k \le K} \big(|C_{J,gf}(k)| + |C_{J,f'}(k)| \big) \le \sum_{k \le K} |C_{J,gf}(k)| + c. \end{aligned} \tag{4.10}$$

The above inequality implies that if $C_{J,\,gf}(k)$ is absolutely summable, the linear response holds. However, the converse is not necessarily true, which means that the summability of $C_{J,gf}(k)$ is a sufficient, but not necessary condition, for the smoothness of statistics. There could exist weaker conditions for $C_{J,gf}(k)$, under which the linear response still exists even beyond the universe of uniformly hyperbolic systems. Assuming both J and f belong to the class of arbitrarily smooth functions,

the summability of $C_{J,gf}(k)$ is guaranteed if g belongs to the space of observables with exponential decay of correlations. For any h in that function space, the Central Limit Theorem (CLT) holds for $h, h\circ\varphi, h\circ\varphi_2, ...$ [194, 109, 36] (see also general large deviation estimates for continuous maps in [193]). In this context, CLT implies that the random variable $\frac{1}{\sqrt{N}}\sum_{n=0}^{N-1}\left(h\circ\varphi_n(x) - \langle h\rangle\right)$ is distributed according to a normal distribution with mean 0 and a finite variance for large N and almost every initial condition x.

Based on the above discussion, we conclude that a comprehensive analysis of the distribution of $|g|$ could provide us further hints into the validity of the linear response. We expect that the statistics is differentiable if CLT holds. This conjecture could be easily verified using Algorithm 2 to compute the time series of g along a typical trajectory and evaluating the second moment of the underlying distribution. Appendix 4.7.1 promptly describes a version of that algorithm adjusted to one-dimensional maps. In the remainder of this chapter, we will leverage our numerical methods for g to verify our analysis. For that purpose, we develop an independent method verifying the differentiability of statistics based. In particular, we approximate the Hölder exponent of one-dimensional to alternatively assess the smoothness of the linear response. Our numerical experiments indicate that the validity of CLT is indeed a sufficient condition for the existence of the linear response. Moreover, it turns out that Lebesgue-integrability of g, i.e, $g \in L^1(\rho)$, could also be sufficient.

4.4 Probing the differentiability of statistics of one-dimensional chaos

The central part of this chapter is to identify a computable mathematical criterion for the differentiability of statistics. As argued in Section 4.3, we anticipate there is a relation between the properties of the distribution of $|g|$ and the validity of the linear response. We intend to further investigate this observation using numerical simulation. We compute the distribution of the absolute value of the density gradient

(using Eq. 4.19) over a range of parameter values. We also compute the statistics-parameter curve directly, from which we estimate its Hölder exponent. Accurate estimation of the Hölder exponent requires a huge amount of samples. We can still afford it given the low dimension of considered examples.

4.4.1 Visualizing the distribution of the ergodic measure gradient

We first study the distribution of $|g|$, considering it to be a random variable, for the onion map introduced in Section 4.2. We empirically obtain the distributions of $|g|$ at different γ values, from both the smooth and non-smooth regions. Based on the procedure introduced in Section 4.7.1, we compute a sufficiently long trajectory, and count the number of occurrences of $|g|$ in all bins, each corresponding to a subset of the range of $|g|$. Figure 4-4 shows the distribution of $|g|$ on a logarithmic scale at fixed $h = 0.97$ for different values of the parameter γ.

The vertical axis of Figure 4-4 represents the number of appearances of a given value of $|g|$ in each bin. Based on these histograms, we conclude that the probability density function (PDF) of $|g|$, denoted as $\mathrm{PDF}(|g|)$, has power-law behavior, i.e.

$$\mathrm{PDF}(|g|) \sim |g|^{-t},$$

for some exponent t. Since Figure 4-4 presents data on a log-log scale, the bin size increases proportionally to the value of $|g|$. Therefore, in Figure 4-4, we observe a distribution that is proportional to

$$|g| \cdot \mathrm{PDF}(|g|) \sim |g|^{-t+1}.$$

That is, the slope of the histograms gives us the value of $-t + 1$.

For power-law probability distributions, the existence of the expected value, variance, and higher-order moments, is solely determined by the value of the exponent t.

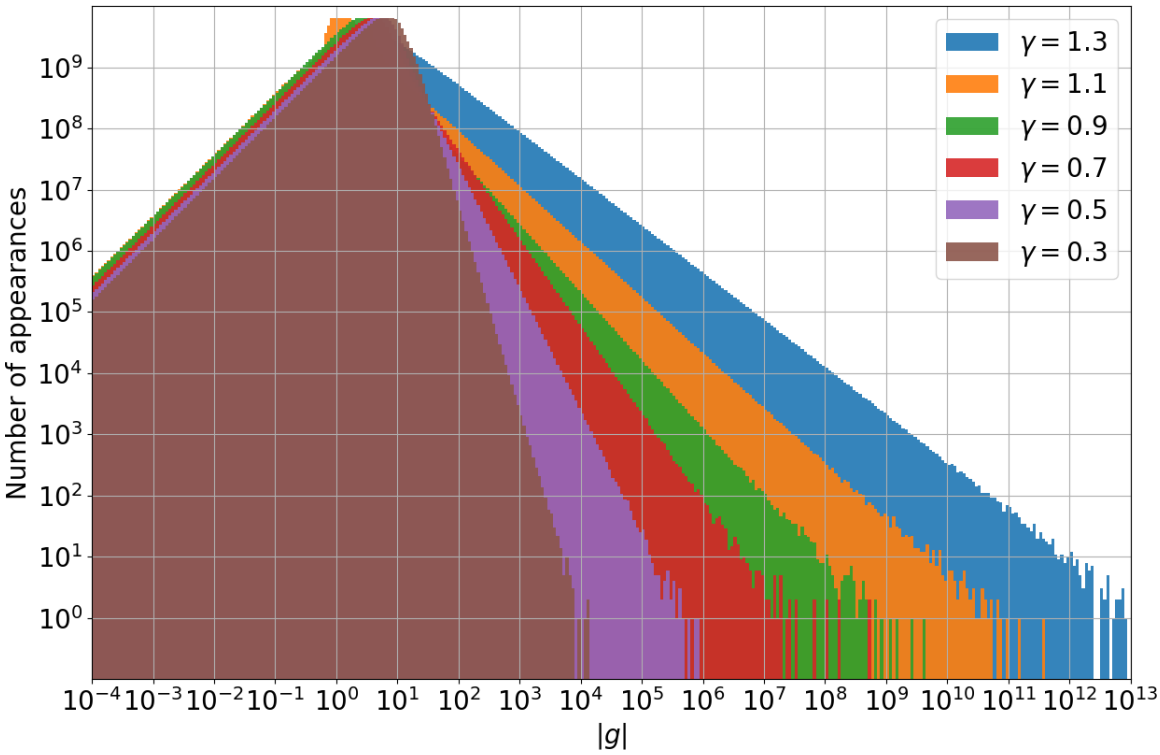

Figure 4-4: Distribution of the absolute value of the density gradient function generated for the onion map (Eq. 4.1) at $h = 0.97$. To generate these histograms, we divided the x-axis from 10^{-18} to 10^{84} into $K = 2048$ bins with equal with in the logarithmic scale. For each histogram, a trajectory of the length of approximately $N = 1.25 \cdot 10^{11}$ has been computed. Along the trajectory, no instances of $|g| > 10^9$ have been identified by the algorithm.

If $t \leq 2$, then the mean (expected value),

$$\mathbb{E}[|g|] = \int_0^\infty |g| \, \mathrm{PDF}(|g|) \, d|g| = \int_0^1 |g(x)| \, \rho(x) \, dx,$$

and all higher moments are infinite. If $t > 2$, then the mean is finite. In other words, the density gradient belongs to $L^1(\rho)$, when the slope of the histogram in Figure 4-4 is less than -1. In addition, if the exponent t is larger than 3, then the variance $\mathrm{var}(|g|) = \mathbb{E}[g^2] - (\mathbb{E}[|g|])^2$ of the probability distribution function is finite. In this case, the density gradient is square-integrable with respect to ρ, i.e., $g \in L^2(\rho)$.

From Figure 4-4, the slope at $\gamma \leq 0.9$ is always less than -1. Thus, the absolute value of the density gradient has finite expected value, and g is therefore Lebesgue-integrable as long as $\gamma \leq 0.9$. Thus, the Lebesgue-integrability threshold must be

in $(0.9, 1.1)$, as all the distributions corresponding to $\gamma \geq 1.1$ have slopes t larger than -1. Furthermore, square-integrabilty threshold can be estimated to be around $\gamma = 0.5$, since the slope at $\gamma = 0.5$ is approximately equal to -2. For smaller values of γ, for example, $\gamma = 0.3$, we see that both expectation and variance are finite.

4.4.2 Hölder exponent test

We need an independent numerical technique to infer a correlation between the existence of moments of $|g|$ and the validity of linear response. For this purpose, we directly assess the smoothness of the statistics illustrated in Figure 4-3 for the onion map. That is, we numerically estimate the Hölder exponent $\mu \in (0, 1]$ of the long-time average function, which changes with γ. A function $h(\gamma) : D \to \mathbb{R}$ is Hölder continuous in D with exponent γ, if there exists a $C > 0$ such that

$$|h(\gamma_1) - h(\gamma_2)| \leq C|\gamma_1 - \gamma_2|^{\mu} \tag{4.11}$$

for all possible pairs of γ_1 and γ_2 belonging to D. If $\mu = 1$, then $h(\gamma)$ is Lipschitz-continuous, and in this case, also differentiable at almost every parameter value in D. Thus, to probe the smoothness of the long-time average, we numerically estimate the Hölder exponent of the statistics-parameter relation, $\langle J \rangle$ vs. γ, from the plot in Figure 4-3. This can be achieved by generating a sufficient number of data points and producing a scatter plot with $|\langle J \rangle(\gamma_1) - \langle J \rangle(\gamma_2)|$ on the y-axis and $|\gamma_1 - \gamma_2|$ on the x-axis, where γ_1 and γ_2 indicate points of evaluation of $\langle J \rangle(\gamma)$. If the logarithmic scaling is used, the Hölder exponent μ can be approximated by estimating the slope (steepness) of the maximum values of $|\langle J \rangle(\gamma_1) - \langle J \rangle(\gamma_2)|$ as $|\gamma_1 - \gamma_2|$ changes. The above procedure is in fact a version of detrended fluctuation analysis (DFA) [136].

Assuming the function $\langle J \rangle$ is sampled every $\delta\gamma$ along the x-axis, it is clear that $\delta\gamma = \min_{\gamma_1, \gamma_2 \in D} |\gamma_1 - \gamma_2|$. We set $\delta\gamma = 0.0001$, which allows us to capture high-frequency oscillations. The value of $\delta\gamma$, however, cannot be too small, as the growing statistical noise may significantly impact the value of $|\langle J \rangle(\gamma_1) - \langle J \rangle(\gamma_2)|$. To further reduce the effect of statistical noise, we run 10 independent simulations per one pa-

rameter value and compute the 3-sigma confidence interval of the data coming from these independent simulations, where the standard deviation is averaged over a chosen interval of γ.

The left-hand side column of Figures 4-5 and 4-6 illustrate the statistics versus parameter dependence at four different intervals of γ. The second column of these two figures shows $|\langle \bar{J} \rangle(\gamma_1) - \langle \bar{J} \rangle(\gamma_2)|$ versus $|\gamma_1 - \gamma_2|$ computed from the data presented in the left-hand side column, where $\langle \bar{J} \rangle$ represents the long-time average of a modified objective function $\bar{J}$. The new quantity of interest is obtained by subtracting a linear function from $\langle J \rangle$, illustrated in Figure 4-3, such that the resulting long-time average vanishes at the end points of each interval of γ. This modification is made to visually amplify the roughness of the curve, which is done for demonstration purposes only. See the caption of Figure 4-5 for more details.

The top row of Figure 4-5 corresponds to the range $\gamma \in [0.2, 0.45]$. It is evident that the long-time average is smooth in this interval, as it satisfies Ineq. 4.11 with the exponent $\mu \approx 1$. According to Figure 4-4, the tail of the distribution of $|g|$ has a slope smaller than -2 in that interval, which implies that $|g|$ has both finite mean and variance. The second row of Figure 4-5 corresponds to the interval $\gamma \in [0.65, 0.9]$, in which the statistics seems sharper, but the Hölder exponent μ is still close to 1. Figure 4-4 indicates that g is in $L^1(\rho)$ but not in $L^2(\rho)$ in that range (it has finite mean, but infinite variance), as the slope of the distribution tail is between -2 and -1. Figure 4-6 includes two sets of plots showing clearly non-smooth, even discontinuous responses. Even for the interval $\gamma \in [1.1, 1.35]$, the Hölder exponent is significantly smaller than 1, which indicates that the statistics are not differentiable. In case of $\gamma \in [1.55, 1.8]$, the long-time average is not even Hölder-continuous with respect to γ. Again correlating with our numerical results on the distribution of $|g|$, for $\gamma > 1$, we found that $|g|$ does not have a finite mean.

The plots in the right column of Figures 4-5 and 4-6 clearly indicate that the statistics of the onion map is differentiable as long as γ is smaller than 1. We refer the reader to the work of Baladi and Smania [12] involving a rigorous study of the regularity of statistics for piecewise expading unimodal maps. Note the onion map,

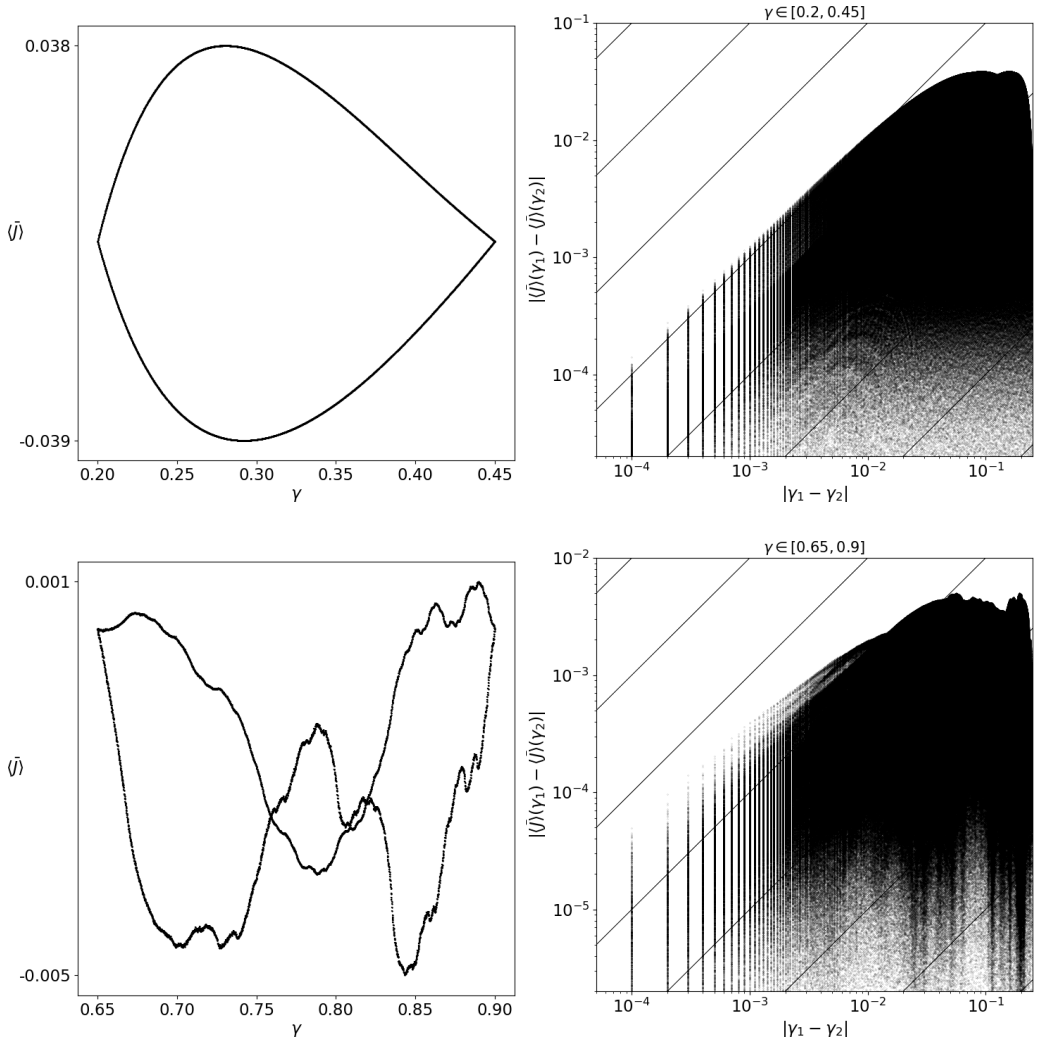

Figure 4-5: Left column: relation of the long-time average and the exponent γ for the onion map at $h = 0.97$. The simulation data is the same as the data presented in Figure 4-3 for $c = 0.625$, however the quantity of interest has been changed to $\langle \bar{J} \rangle$ by subtracting a linear function describing a straight line crossing the endpoints of the curve in Figure 4-3 in each γ interval from the original objective function. Each plot corresponds to a different γ interval between γ_{min} and γ_{max}, which has been discretized uniformly with step size $\delta\gamma = 0.0001$. For each value of γ, we run 10 independent simulations. Right column: Hölder exponent test results. First, for each pair of data points from the left-hand side plot, excluding the pairs with the same value of γ (i.e. when $\gamma_1 = \gamma_2$), we compute the difference of the corresponding long-time average values versus the difference of their parameter values. Second, we compute the lower-bound of the 3-sigma confidence interval by subtracting 6 averaged sigmas, where sigma represents standard deviation of results obtained in 10 simulations averaged over the interval $[\gamma_{min}, \gamma_{max}]$, from the computed differences of modified statistics. This means each plot has approximately $(0.5 \cdot (\gamma_{max} - \gamma_{min})/\delta\gamma)^2 \approx 1.6 \cdot 10^6$ data points. Skew solid lines represent reference lines with the slope of 1 in the logarithmic scale.

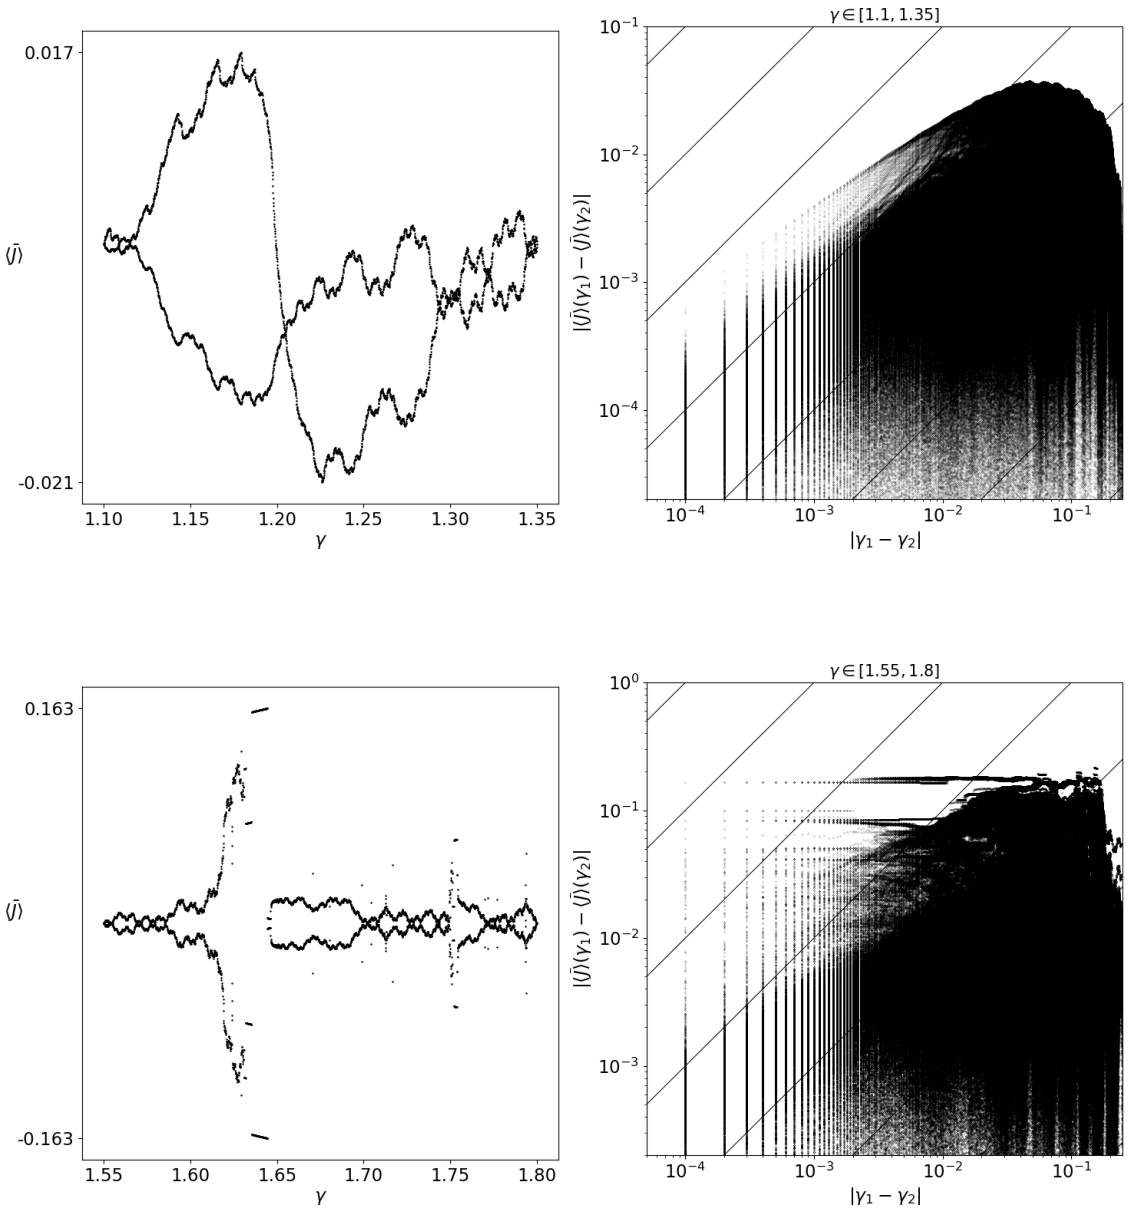

Figure 4-6: This figure is an extension of Figure 4-5. It includes γ-intervals corresponding to non-smooth statistics. All plots have been generated in the same manner as those in Figure 4-5 – see the corresponding caption for more details.

at $h = 0.97$, is a piecewise expanding unimodal map for all values of $\gamma \in [0.24, 0.99]$ (the boundary values are approximated up to 0.005). According to our analysis of the distribution of $|g|$, $\gamma < 1.0$ implies finite expectation of $|g|$. It means that, in case of the onion map, linear response holds when $g \in L^1(\rho)$. This result also confirms our analysis from Section 4.3. From our numerical results, in this case, we find that the converse is also true: when $g \notin L^1(\rho)$, linear response fails. To check whether the equivalence

$$g \in L^1(\rho) \Longleftrightarrow \left| \frac{d\langle J \rangle}{d\gamma} \right| < \infty$$

is generalizable, we will apply the above two-step procedure to a higher-dimensional system with one positive LE.

4.5 Generalization to multi-dimensional flows: Example of Lorenz 63

In this section, we generalize our conclusions from Section 4.4 to higher-dimensional systems with a one-dimensional unstable manifold. This means we consider n-dimensional systems that have exactly one positive Lyapunov exponent out of n Lyapunov exponents. As a test case, we consider the Lorenz 63 system [110, 166], which consists of three coupled nonlinear ODEs,

$$\frac{dx^{(1)}}{dt} = \sigma(x^{(2)} - x^{(1)}), \quad \frac{dx^{(2)}}{dt} = x^{(1)}(\gamma - x^{(3)}) - x^{(2)}, \quad \frac{dx^{(3)}}{dt} = x^{(1)}x^{(2)} - \beta x^{(3)}, \tag{4.12}$$

where $\sigma \geq 0$, $\beta \geq 0$, and $\gamma \geq 0$ are the system parameters. This ODE system models thermal convection of a fluid cell that is warmed from one side and cooled from the opposite side. The solution to Eq. 4.12 is represented by a 3-element state vector $x(t) = [x^{(1)}(t), x^{(2)}(t), x^{(3)}(t)]^T$. In our analysis, we set σ and β to their canonical values of 10 and $8/3$, respectively, and keep them fixed, while we allow γ to vary. Given the Lorenz '63 system is a three-dimensional system, it has three distinct Lyapunov exponents λ_i, $i = 1, 2, 3$, indexed in decreasing order. They satisfy the

following constraints [166],

$$\lambda_1 + \lambda_2 + \lambda_3 = -(1 + \sigma + \beta), \quad \lambda_2 = 0. \tag{4.13}$$

Since both parameters are assumed to be positive, it is evident that Eq. 4.13 admits at most one positive solution. According to [166], for the canonical values of σ and β, the Lorenz 63 system is:

- non-chaotic (has no positive LEs) if $0 \leq \gamma < 24.7$ and $\gamma > 99.5$,
- chaotic (has one positive LE) if $24.7 \leq \gamma \leq 99.5$.

Therefore, in this section, we focus on the smoothness of statistics of the Lorenz 63 system when $\gamma \in [24.7, 99.5]$. Several computational studies have found that the response of this system to small parameter perturbations at $\gamma = 28$ (canonical value) is linear, i.e., long-time averages of smooth obervables are differentiable at $\gamma = 28$ [72, 144]. To generate all results presented in this section, we integrate the system forward in time using the second-order Runge-Kutta scheme (midpoint method) with time step[2] Δt, starting from a random initial vector x_{init}. In our discussion, we no longer consider the original, that is, continuous version of Lorenz 63, but rather we focus on the discrete form using a map φ, which is defined by the numerical time integration of the Lorenz 63 system for a time of Δt; that is,

$$\varphi(x(t)) = x(t + \Delta t)$$

for all $t \in \mathbb{R}^+$. In other words, a discretized orbit of φ, represented by a sequence of 3-dimensional state vectors, x_0, x_1,..., $x_k := x(k\,\Delta t)$, $k \in \mathbb{Z}^+$, is a numerical solution of the Lorenz equations. To make sure the numerical integration error is within a desired tolerance, one could also consider an adaptive time-stepping scheme [165]. While an extra computational effort would be required to compute Δt every time step, the approximated trajectory might be closer to the true solution.

[2]Specific values of the time step size Δt are indicated in the captions of corresponding figures.

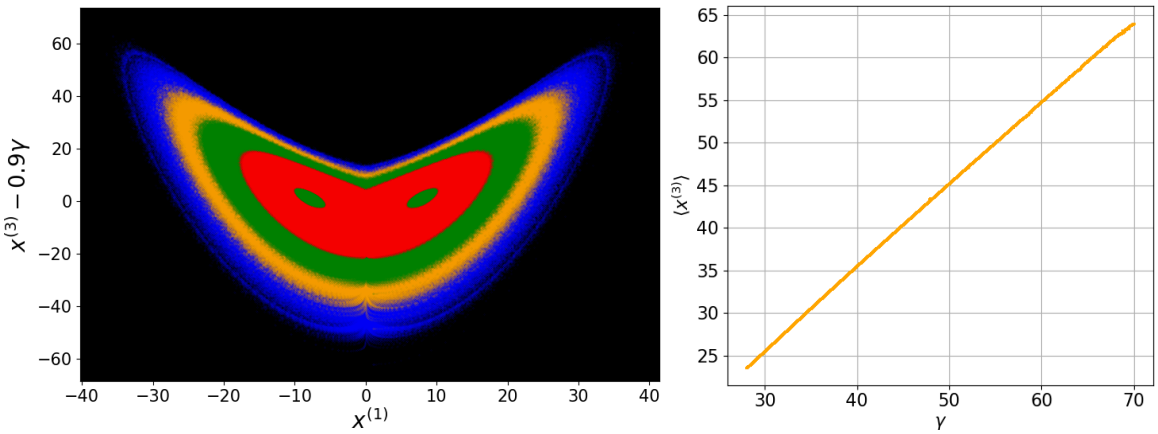

Figure 4-7: Left: projection on the $x^{(1)}$-$x^{(3)}$ plane of the Lorenz 63 attractor at $\gamma = 25$ (red), 40 (green), 55 (orange) and 70 (blue). Each projection has been shifted downwards proportionally to γ for demonstration purposes. Right: relation of the long-time average $\langle x^{(3)} \rangle$ (defined by Eq. 4.14) and the system parameter γ. We generate 420,000 data points in total: for each value of γ on a uniform grid with size $\delta\gamma = 0.001$, we run 10 independent simulations. For each data point, we compute approximately $2.5 \cdot 10^{10}$ time steps with $\Delta t = 0.01$.

As a quantity of interest, we consider the long-time average of the third variable,

$$\langle x^{(3)} \rangle = \lim_{N \to \infty} \frac{1}{N} \sum_{k=0}^{N-1} x_k^{(3)}, \tag{4.14}$$

which we approximate as $\langle x^{(3)} \rangle$ by generating sufficiently long trajectories. Figure 4-7 shows a 2D projection of the attractor at different values of γ, as well as the dependence of $\langle x^{(3)} \rangle$ on γ. We observe that the attractor expands outward on the $x^{(1)}$-$x^{(3)}$ plane, as γ increases. This observation is also reflected in the linear relation between $\langle x^{(3)} \rangle$ and γ, which is shown on the right-hand side of Figure 4-7. To show the statistical quantities of the Lorenz 63 system are in fact non-smooth at some values of γ, we subtract a smooth function $s(\gamma)$, obtained by fitting $x^{(3)}$ vs. γ with a quadratic polynomial, from the original data shown in the right plot of Figure 4-7. Figure 4-8 illustrates the behavior of the modified quantity of interest, i.e., $x^{(3)} - s(\gamma)$, as γ changes. This computational treatment clearly reveals the actual regularity of the system's statistics. Analogously to the one-dimensional onion map, here as well we observe a transition from a smooth response at lower values of γ to a non-smooth,

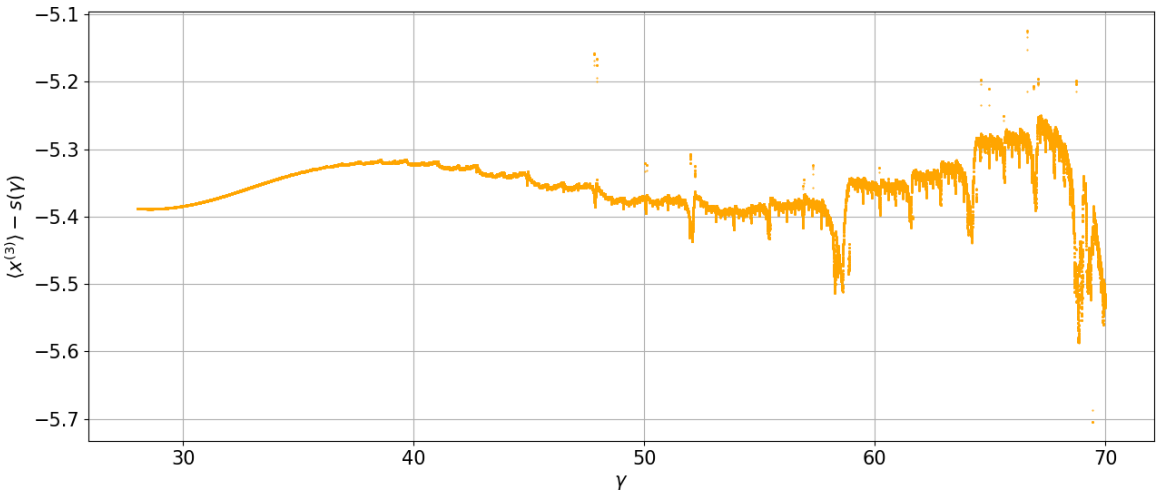

Figure 4-8: The modified quantity of interest, $\langle x^{(3)} \rangle - s(\gamma)$, as a function of γ, where $s(\gamma) = 1.06\gamma - 0.00095\gamma^2$. The quadratic function $s(\gamma)$ was obtained using a two-step procedure. First, as a starting point, we applied least squares polynomial fitting to the original long-time average shown in Figure 4-7. Subsequently, we manually adjusted the coefficient proportional to the linear term in order to expose the rough behavior of the curve. Recall the modification of the quantity of interest was made for demonstration purposes only.

and even discontinuous, behavior at larger values of γ. Note the modified statistics becomes sharp for values of γ slightly above 30. This is in fact the region where the Lorenz 63 system loses its quasi-hyperbolic properties [166].

For completeness, we illustrate the $x^{(1)}$–$x^{(3)}$ projection of the empirical density function of the Lorenz 63 system at three different values of γ in Figure 4-9. We notice a clearly smooth distribution for $\gamma = 28$. For $\gamma = 38$, however, subtle wrinkles are visible around the "eyes" of the attractor. In case of $\gamma = 70$, regions with large density gradients, which clearly indicate non-smoothness of the distribution, appear around the "eyes" and close to the boundary of the attractor.

Based on these observations, we anticipate the density gradient function to be smooth for values of γ close to 28, and non-smooth if γ is higher. We also acknowledge a consistency between Figures 4-8–4-9 and Figures 4-2–4-3, corresponding to the Lorenz 63 system and onion map, respectively. Both pairs of figures indicate a strong correlation between the smoothness of statistics and smoothness of the density function in phase space. To infer a more rigorous connections, we will apply the

numerical methodology described in Section 4.4 for Lorenz 63.

4.5.1 Uniform hyperbolicity and Ruelle's formula of flows with one positive LE

We argue that the conclusion inferred in Section 4.3 is applicable to multi-dimensional systems with one-dimensional unstable manifolds, such as Lorenz 63. We show that a part of the full response is structurally similar to the regularized expression derived for one-dimensional maps (Eq. 4.5).

Let us consider an invertible, ergodic, discrete map of a manifold M, parameterized by γ, and given by

$$x_{k+1} = \varphi(x_k; \gamma), \quad k \in \mathbb{Z}. \tag{4.15}$$

Here, x_k is an n-dimensional state vector. Eq. 4.15 may arise from the time discretization of an ODE system governed by the evolution vector h. Let D denote the phase-space gradient operator with $D\varphi$ being the $n \times n$ Jacobian matrix of the system. For the system defined by Eq. 4.15, Ruelle's formula [148, 150] for the parametric derivative of the long-time average can be expressed as

$$\frac{d\langle J \rangle}{d\gamma} = \sum_{k=0}^{\infty} \int_M D(J \circ \varphi_k) \cdot \chi \, d\mu, \tag{4.16}$$

where $\chi := \partial\varphi/\partial\gamma \circ \varphi^{-1}$ is the parametric perturbation vector, while μ is the underlying SRB measure. Eq. 4.16 is rigorously true if φ is a uniformly hyperbolic system. The tangent space of such systems admits the following decomposition,

$$T_x M = E^u(x) \oplus E^s(x) \oplus E^c(x),$$

which satisfies the following properties:

- covariance property:

$$D\varphi(E^u(x)) = E^u(\varphi(x)), \; D\varphi(E^s(x)) = E^s(\varphi(x)), \; D\varphi(E^c(x)) = E^c(\varphi(x)),$$

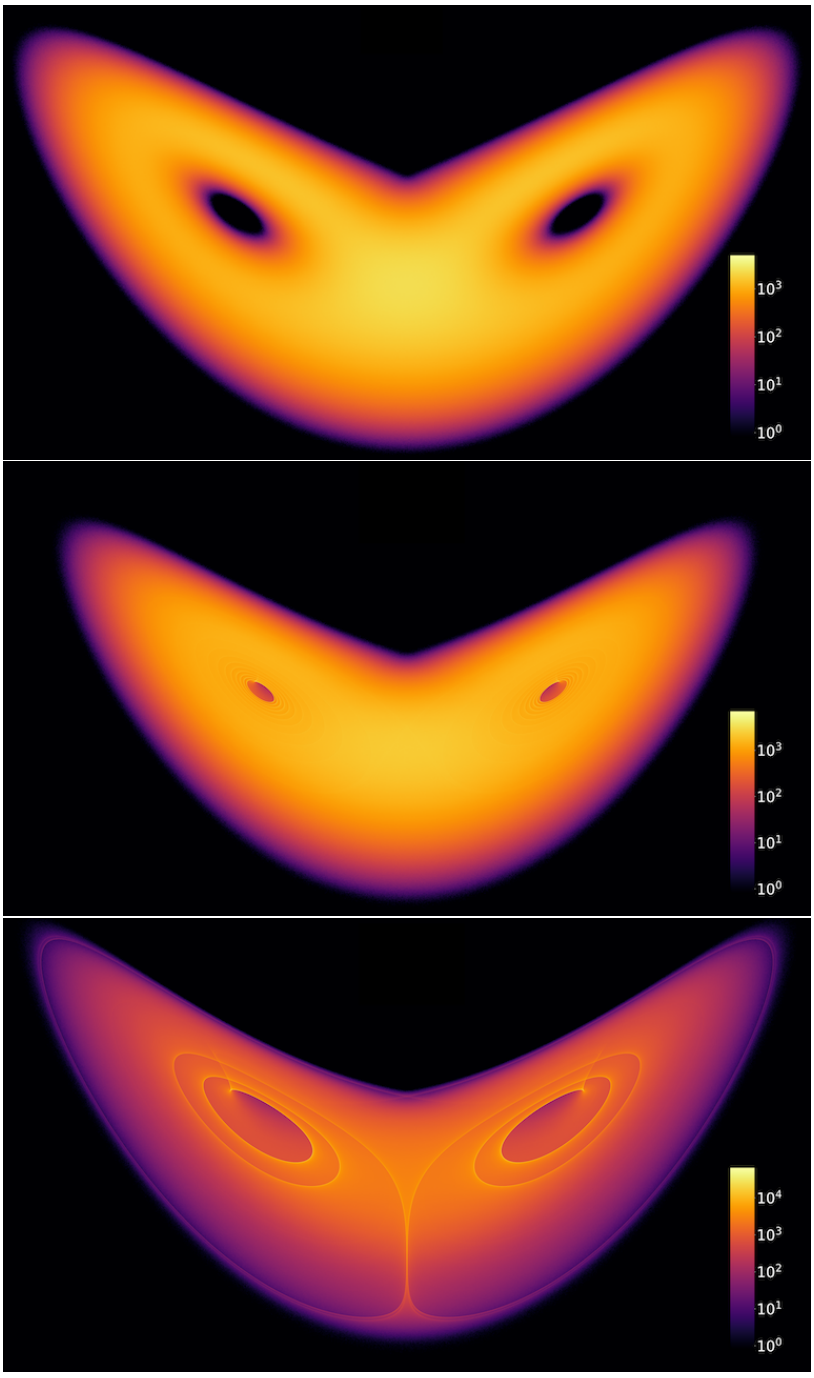

Figure 4-9: $x^{(1)} - x^{(3)}$ projection of the (unnormalized) empirical density function of the Lorenz 63 system at $\gamma = 28$ (top), 38 (middle) and 70 (bottom). To generate each plot, a 2D box in phase space has been divided into $n_{x^1} \cdot n_{x^3} = 3840 \cdot 2160$ uniform rectangular cells/bins. The coordinates of the bottom left/upper right corner of each box are the following: $[-20, 0]/[20, 50]$ (top), $[-27, 0]/[54, 70]$ (middle), $[-40, 0]/[40, 80]$ (bottom). We computed a trajectory of length $N = 0.5 \cdot 10^9$ with $\Delta t = 0.002$ for each plot. The color bars indicate the number of times the trajectory crosses a bin.

- uniform expansion/contraction:
  for some fixed constants $C > 0$, $\lambda \in (0, 1)$, at every $x \in M$, every vector $v \in E^u(x)$ satisfies
  $$\|D\varphi_{-k}(x)\ v(x)\| \leq C\lambda^k \|v(x)\|$$
  for all positive integers k. And,
  $$\|D\varphi_k(x)\ v(x)\| \leq C\lambda^k \|v(x)\|$$
  for all $v \in E^s(x)$. The norm, $\|\cdot\|$ denotes the standard Euclidean norm in $\mathbb{R}^n$,

- neutral subspace E^c is one-dimensional and aligned with φ (parallel to the evolution vector of the continuous system):
  at every $x \in M$, there exist constants $C \in \mathbb{R}$ and $C_1 \in \mathbb{R}$ such that
  $$v = C\, h(x) \approx C_1(\varphi(x) - x)$$
  for all $v \in E^c(x)$.

If the above properties hold, the input perturbation χ can be linearly decomposed as $\chi_1 + \chi_2 + \chi_3$ such that

- $\chi_1(x) \in E^u(x)$, $\chi_2(x) \in E^c(x)$, and all the three components are differentiable on the unstable manifold;

- there exists a bounded vector field $v : M \to \mathbb{R}^d$ that is orthogonal to $E^u(x) \oplus E^c(x)$

The part of the linear response involving χ_3 will later be referred to as the *stable* contribution. It can be rigorously re-expressed through a single ergodic average of $DJ \cdot v$ such that v is an inhomogeneous tangent that is orthogonal to the center-unstable manifold everywhere on the attractor. Assuming $\|DJ\|_\infty < \infty$, the stable contribution is guaranteed to be bounded. The sensitivity due to χ_2, i.e., the *neutral* or *center* contribution is bounded as well, because it can be expressed in terms of

the time-correlation between a bounded scalar and directional derivative of J in the direction parallel to the flow, both of which are assumed to be Hölder-continuous. The reader is referred to [32] and Chapter 6 for more details regarding these two terms.

We now consider the last remaining term of the splitting, χ_1, which we call the *unstable* contribution. If the system has only one positive LE, χ_1, is always tangent to the unstable manifold, geometrically represented by a curve. Let $q(x)$ be a unit vector parallel to the unstable manifold at x and $\chi_1 = c\,q$. Therefore, the unstable contribution equals

$$\sum_{n=0}^{\infty} \int_M D(J \circ \varphi_n) \cdot \chi_1 \, d\mu = -\sum_{n=0}^{\infty} \int_M (J \circ \varphi_n)\,(cg + b)\, d\mu, \tag{4.17}$$

where b denotes the directional derivative of c in the q direction, while g is the SRB measure gradient, as defined by Eq 3.62. The reader is referred to Chapters 5-6 for a detailed derivation of Eq. 4.17 and its generalization to systems with more than one LE. Note that if M itself is a one-dimensional manifold, as in the onion map, the unstable contribution, given by Eq. 4.17 is the entire sensitivity, since there are no stable nor center contributions. We highlight the fact that Eq. 4.17 is almost identical to its one-dimensional counterpart (Eq. 4.5). Thus, the connection between the regularity of g and the existence of linear response that we described in Section 4.3 can potentially be extended to multi-dimensional systems with a one-dimensional unstable manifold (e.g., the Lorenz 63 system). Later, in Chapter 5, it will be clear that this methodology is easily extendable to systems with high-dimensional unstable manifolds.

We acknowledge that the Lorenz 63 oscillator is a partially hyperbolic system at certain values of γ [24]. Partial hyperbolicity means that, in the splitting $T_xM = E^u \oplus E^s \oplus E^c$, the stable and unstable subspaces satisfy the uniform contraction/expansion condition and the dominance condition with respect to the central subspace E^c (see [24] for a rigorous definition of these conditions). In the work of Dolgopyat [46], the linear response theory has been rigorously extended to partially hyperbolic sys-

tems. Therefore, the presented logic using the definition of uniform hyperbolicity is applicable to Lorenz 63 as well.

4.5.2 Probing the differentiability of statistics of Lorenz 63

Using Algorithm 1, we compute g and generate histograms of the absolute value of the density gradient function $|g| = |\partial_\xi \log \rho|$ for the Lorenz 63 system. Figure 4-10 illustrates the distributions of $|g|$ at three different values of γ. We observe power-law behavior of the generated histograms similar to those of the onion map in Figure 4-4. Clearly, the exponent t, which is introduced in Section 4.4, is much higher than 3 if $\gamma = 28$. This implies that both the expected value and variance of $|g|$ are finite, which means g is square-integrable (with respect to ρ). The other two distributions (at $\gamma = 40$ and $\gamma = 68$) feature tails with exponents t slightly smaller than 2, which means that g may not be Lebesgue-integrable, as discussed in Section 4.4.

Figure 4-10 clearly indicates that the Lebesgue-integrability threshold can be estimated to be at some γ between 28 and 40. This result can be correlated with the regularity of the density function ρ (see Figure 4-9), which apparently loses its global smoothness for $\gamma \leq 70$. In an extensive study of the Lorenz 63 attractor at canonical values of β and σ presented in [166], it was shown that the system is quasi-hyperbolic if $\gamma \in [24.06, 31]$ and non-hyperbolic if $\gamma \in [31, 99.5]$. Quasi-hyperbolicity means that some critical features of uniformly hyperbolic systems, e.g., smoothness of the physical measure describing the statistical behavior on the attractor, still apply [18]. Thus, our results confirm that, in this case, the loss of (quasi-)hyperbolicity is also an indicator of the failure of linear response.

We now estimate the Hölder exponent μ, as defined in Eq. 4.11, for the statistics-vs-parameter relation presented in Figure 4-8, using the procedure described in Section 4.4.2. Figures 4-11-4-12 illustrate the results of the Hölder exponent numerical test generated for three different intervals of γ. These results clearly indicate the exponent μ is approximately 1 if $\gamma \in [28, 32]$, implying Lipschitz-continuity of that part of the curve. The plots in the bottom row of Figure 4-11 and Figure 4-12 show μ is significantly smaller than 1, which implies the long-time average cannot be dif-

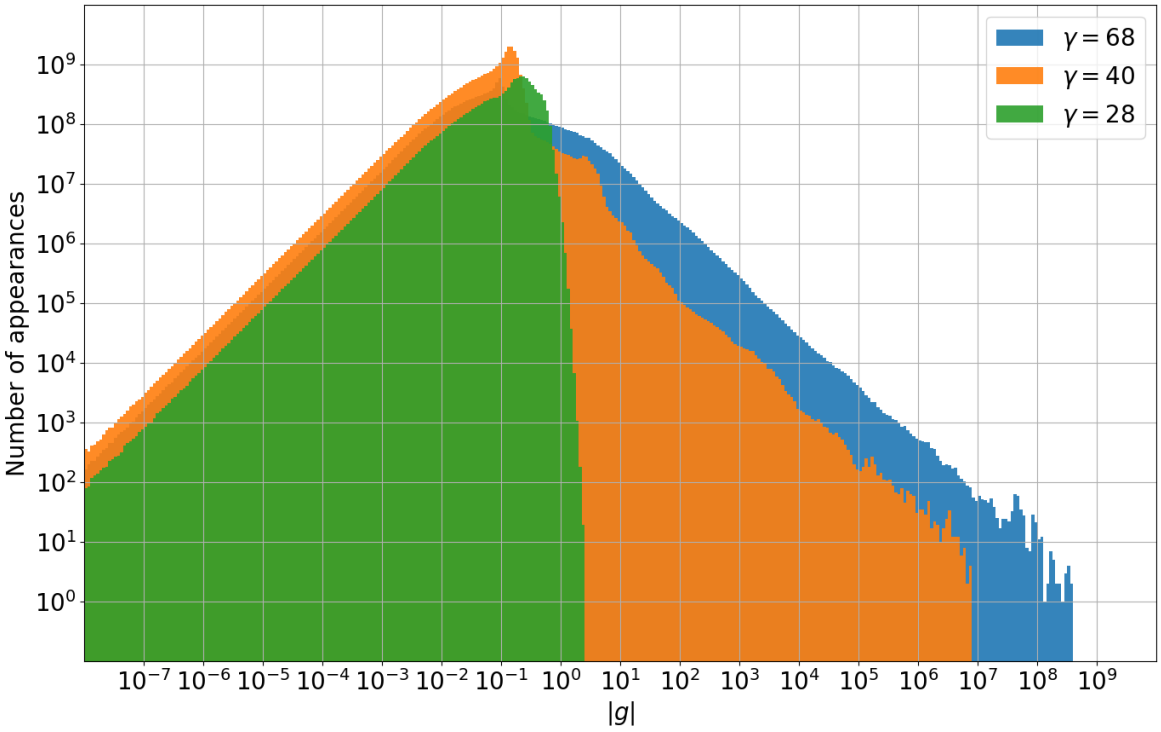

Figure 4-10: Distribution of the absolute value of the density gradient function generated for the Lorenz 63 attractor at three different values of γ using Algorithm 1. To generate these histograms, we divided the x-axis from 10^{-18} to 10^{84} into $K = 2048$ bins of equal width in the logarithmic scale. For each histogram, a trajectory of the length of approximately $N = 2.5 \cdot 10^{10}$, computed by solving Eq. 4.12 with $\Delta t = 0.01$, is used. Along the trajectory, no instances of $|g| > 10^9$ have been identified by the algorithm.

ferentiable at $\gamma > 36$. Therefore, one can observe a clear correlation between the Lebesgue-integrability (with respect to ρ) of g and smoothness of statistics, which is consistent with our numerical results of the onion map from Section 4.4.

4.6 Summary

Statistical quantities are critical both in understanding and in applications of chaotic phenomena, such as turbulent flows. In many chaotic dynamical systems, the relation between statistical quantities and system parameters is not smooth. In this chapter, we show that the existence of the parametric derivative of a statistics or long-time average (sensitivity) depends on whether the SRB measure gradient is Lebesgue-

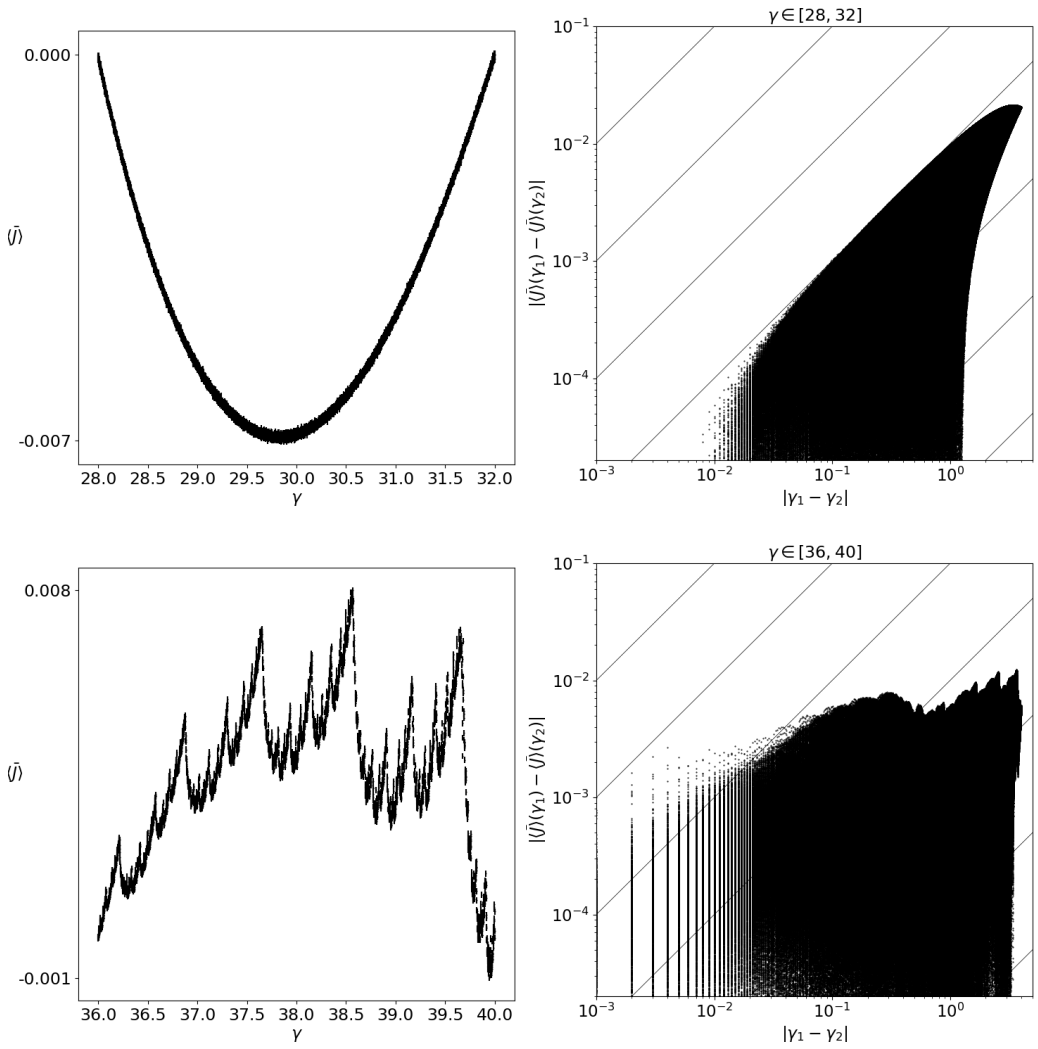

Figure 4-11: Left column: analogously to Figures 4-5 - 4-6, the simulation data is the same as the data presented in Figure 4-8, and the quantity of interest has been modified such that the values of the long-time average at the endpoints of each interval is zero. It has been achieved by subtracting a linear function describing a straight line crossing the endpoints of the original curve, $\langle x^3 \rangle - s(\gamma)$. The modified objective function has been denoted by $\bar{J}$. Each plot corresponds to a different γ interval, which has been discretized uniformly between γ_{min} and γ_{max} with step size $\delta\gamma = 0.001$. For each value of γ, we run 10 simulations. Right column: Hölder exponent test of the statistical quantity $\langle \bar{J} \rangle = \langle z \rangle - s(\gamma)$ versus parameter γ relation plotted in Figure 4-8. These plots have been generated in the same fashion as those for the onion map in Figures 4-5 - 4-6 (see the caption of Figure 4-5 for a detailed description), i.e., by taking the lower bound of the 3-sigma confidence interval of the data set corresponding to $[\gamma_{min}, \gamma_{max}]$, obtained in 10 independent simulations. We sample the statistics every $\delta\gamma = 0.001$, which means each plot has approximately $0.5 \cdot ((\gamma_{max} - \gamma_{min})/\delta\gamma)^2 = 8 \cdot 10^6$ data points. Skew solid lines represent reference lines with the slope of 1 in the logarithmic scale.

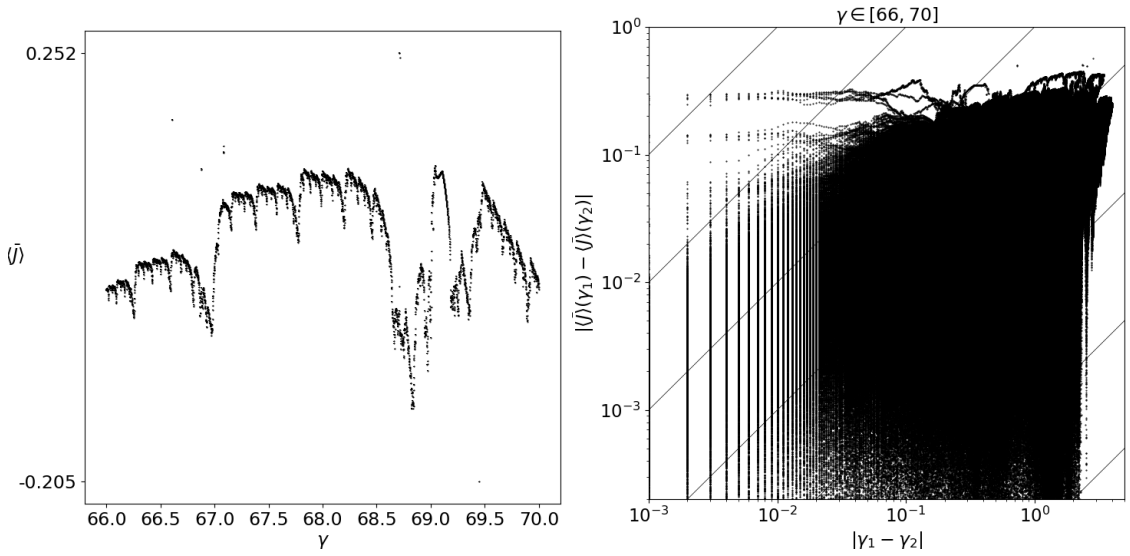

Figure 4-12: This figure is an extension of Figure 4-11. All plots have been generated in the same manner as their counterparts in Figure 4-11 – see caption therein for more details.

integrable. That function represents the relative rate of change of the conditional SRB density with respect to the coordinates of the unstable manifold. The relationship between the sensitivity and g is clearly reflected by the regularized Ruelle's formula, which stems from the linear response theory. This observation can be utilized to construct a simple and generalizable numerical procedure to assess the differentiability of statistics. The computation of the probability distribution of $|g|$ is the central part of the proposed procedure. The probability density function of $|g|$ features a power-law behavior in case of the two systems considered in this chapter: the onion map and Lorenz 63 system. In these special cases, a numerical estimate of the power law exponent is sufficient to determine the differentiability of statistics. We validate this test by numerically computing the Hölder exponent of the one-dimensional output-input relation.

4.7 Appendices

This section includes supplementary material supporting certain arguments presented in this chapter.

4.7.1 Validating the recursion for SRB measure gradient using Ulam's method

The purpose of this section is to present a simple numerical tool based on Ulam's method to verify the iterative formula for g in 1D, which was derived in Section 2.6.3. That formula is a direct consequence of the fact that the SRB density ρ is stationary in time. Indeed, it satisfies

$$\rho(\varphi(x)) = \frac{\rho(x)}{|\varphi'(x)|}, \tag{4.18}$$

which means that ρ is in fact an eigenfunction of the Frobenius-Perron operator with eigenvalue 1. Taking the logarithm of Eq. 4.18 and then differentiating with respect to x, we obtain

$$g(\varphi(x)) = \frac{g(x)}{\varphi'(x)} - \frac{\varphi''(x)}{\varphi'(x)^2}. \tag{4.19}$$

While the measure preservation property in the form of Eq. 4.18 applies only to one-to-one (injective) maps, we showed in 3.10.1 that the resulting recursive formula for g in Eq. 4.19 is in fact applicable to non-injective maps such as the onion map. We highlight that the operator-based approach is another way of deriving the evolution of g.

Eq. 4.19 converges to the true value $\rho'(x)/\rho(x)$ upon iterating with an initial guess $g \circ \varphi_{-N}(x) = 0$, as $N \to \infty$. A convenient way to numerically verify Eq. 4.19 is to approximate $\rho' = \rho g$ using the above formula for g. We validate the results against the finite difference approximation of ρ'. Let $x_0, x_1, ...,$ be a long trajectory, and let the interval $[0, 1]$ be divided into K subintervals (bins), $\{\Delta_k\}_{k=1}^{K}$, of equal length $1/K$. We compute a piecewise-constant approximation of $\rho(x)g(x)$ as follows (analogously to Ulam's approximation [174]; see also [107] for a numerical example involving a simple 1D map):

$$\rho(x)g(x) \approx \frac{K}{N} \sum_{n=0}^{N-1} g(x_n) I_{\Delta_k}(x_n), \; \forall \, x \in \Delta_k \tag{4.20}$$

where I_A is the indicator function over a subset $A \subset [0, 1]$. That is, $I_A(x) = 1$,

when $x \in A$, and $I_A(x) = 0$, otherwise. The pointwise error associated with this approximation is proportional to $\sqrt{K/N}$, i.e.,

$$|\rho(x)g(x) - (K/N)\sum_{n=0}^{N-1} g(x_n)I_{\Delta_k}(x_n)|$$

decays as $\mathcal{O}(\sqrt{K/N})$ for all $x \in \Delta_k$ if g obeys the CLT [36]. Note that, for a fixed N, the error increases proportionally to $\sqrt{K}$ for all $x \in [0, 1]$, because the approximation is piecewise-constant on a uniform grid of size $1/K$. Note if we replace $g(x_n)$ with 1 in the RHS of Eq. 4.20, we effectively obtain a formula for the density function itself. Thus, from the algorithmic point of view, the process of generating ρ' requires similar steps as the process of generating ρ, while g emerges as a byproduct. Analogously, this process can be generalized to higher-dimensional systems with a 1D unstable manifold. In such systems, g is a scalar function, and thus Eq. 4.20 still applies assuming an analogous partition of the higher-dimensional attractor is created.

Figure 4-13 illustrates the derivative of density generated for the onion map for the same set of parameter values as the densities in Figure 4-2. We observe a satisfactory match between the results generated using the above algorithm for $\rho'(x)$ and the corresponding finite difference approximations as long as $\gamma < 1.0$. For larger values of γ, there is a visible discrepancy between the two approximations in the proximity of discontinuities, which is consistent with the density ρ exhibiting discontinuities for $\gamma > 1.0$ (compare with Figure 4-2). We also notice there are significant systematic differences, in terms of the function behavior, between the three parameter regimes: $\gamma < 1$, $\gamma = 1$, $\gamma > 1$. In 4.7.2, we discuss the impact of γ on the formation of singularities due to the onion map transformation.

4.7.2 Singularity formation due to the onion map transformation

The purpose of this appendix is to discuss the behavior of the probability density function (PDF) and the corresponding density gradient resulting from a single onion

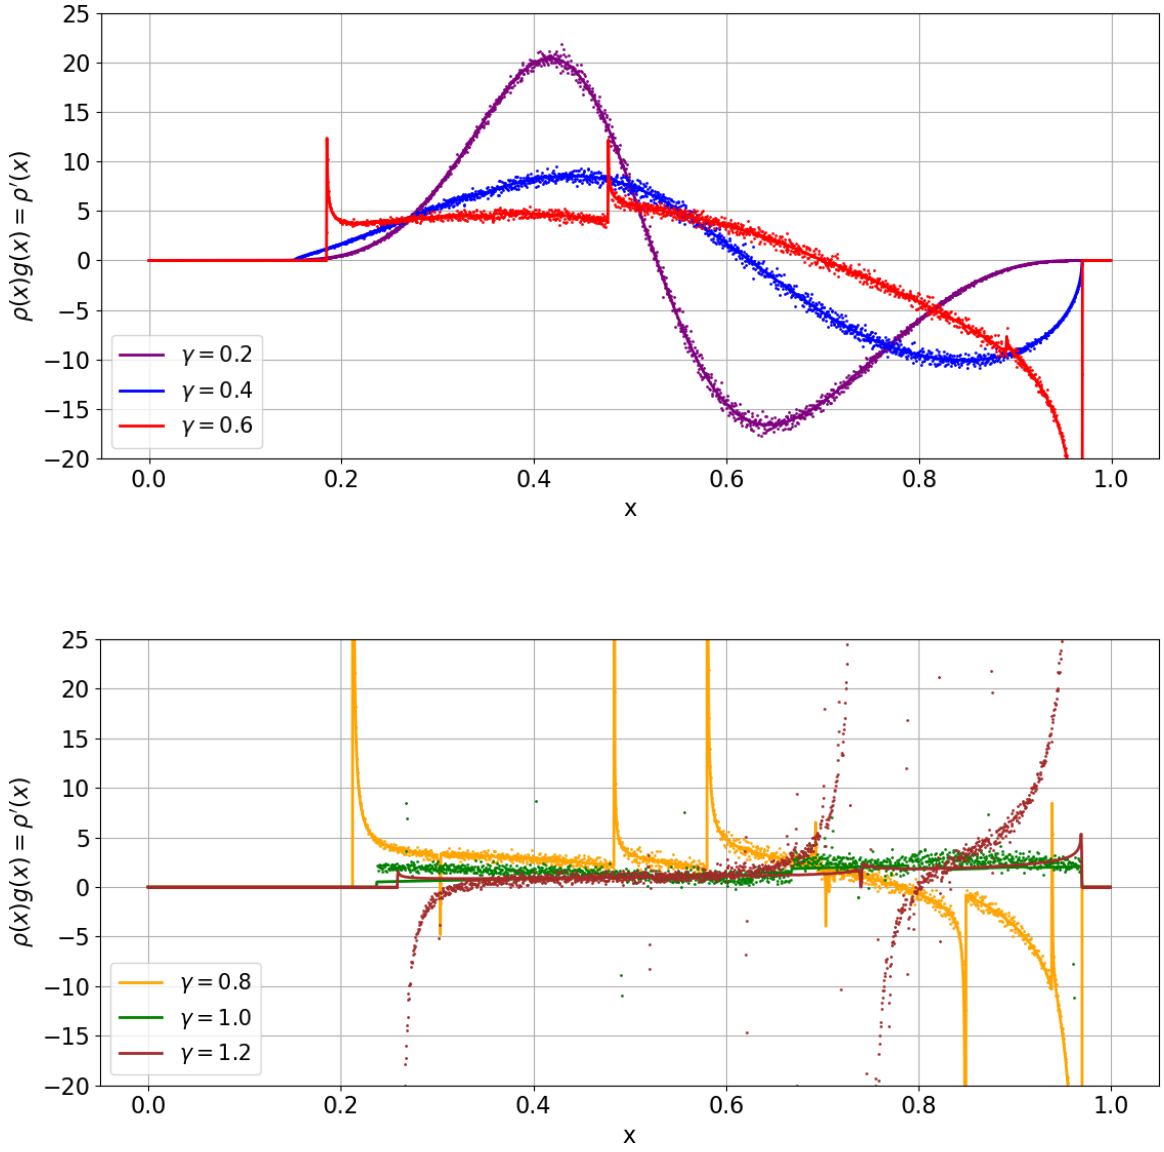

Figure 4-13: Derivative of density of the onion map (Eq. 4.1) at $h = 0.97$. We used $N = 41,943,040,000$ samples and $K = 2048$ bins to generate all curves. The solid lines represent the derivative of density computed using Eq. 4.19-4.20, while the dots represent central finite difference approximation of the same function using the corresponding density function histograms illustrated in Figure 4-2.

map transformation φ applied to uniformly distributed data. This analysis sheds light on the smooth-rough regime transition at $\gamma = 1$. In our derivation, we use generic constants, C_1, C_2, which represent real numbers that do not depend on map parameters.

Let $X \in [0, 1]$ be a uniformly distributed random variable and $Y = \varphi(X)$, where $\varphi : [0, 1] \to [0, h]$ is the onion map defined by Eq. 4.1. The CDF F_Y of Y is thus proportional to

$$\exp(\log(1 - (y/h)^2)/\gamma) = (1 - (y/h)^2)^{1/\gamma}.$$

The first derivative of the CDF equals the PDF f_Y, which can be expressed as follows,

$$f_Y(y) = F_Y'(y) = C_1 \frac{y \left(1 - \left(\frac{y}{h}\right)^2\right)^{\frac{1-\gamma}{\gamma}}}{\gamma h^2}. \tag{4.21}$$

We observe that, after a single transformation, the PDF develops a singularity at $y = h$ if $\gamma > 1$. Note also f_Y is quadratic if $\gamma < 1$, and linear if $\gamma = 1$. The behavior of the SRB density ρ is more complicated than f_Y, since the former reflects the statistics obtained after infinitely many applications of φ. Nevertheless, Figure 4-2, which illustrates the empirical density at different values of γ, is fully consistent with our results. Note if $\gamma > 1$, the SRB density distribution features a bunch of spikes, while at $\gamma = 1$ the distribution seems to be piecewise linear.

Let us now compute the second derivative of the CDF,

$$F_Y''(y) = C_2 \frac{\left(1 - \left(\frac{y}{h}\right)^2\right)^{\frac{1-2\gamma}{\gamma}} (\gamma(h^2 + y^2) - 2y^2)}{\gamma^2 \, h^4}. \tag{4.22}$$

Let the probability density gradient g_Y be defined as the derivative of $\log(f_Y)$, analogously to g. Thus, the expected value of g_Y equals

$$\int_{-\infty}^{\infty} g_Y \, \mathrm{PDF}(g_Y) \, dg_Y = \int_0^h g_Y(y) \, f_Y(y) \, dy = \int_0^h F''(y) \, dy. \tag{4.23}$$

Combining Eq. 4.22 and Eq. 4.23, it is possible to analytically compute the expected

value of g_Y,

$$\int_0^h F''(y)\,dy = \frac{C_2 h^{-\frac{2}{\gamma}}}{\gamma} \lim_{y\to h} \left((h-y)(h+y)\right)^{\frac{1-\gamma}{\gamma}}. \tag{4.24}$$

Since both h and γ are positive numbers, there are two possible scenarios. Namely, the expected value of g_Y is finite (equals zero) if $\gamma \leq 1$, regardless of the value of h. Otherwise, the mean of g_Y does not exist. In any case, the existence of the expected value strictly depends on the behavior around the only spike of $\sim (y-h)^{(1-\gamma)/\gamma}$ at $y = h$. This result also applies to the absolute value of g_Y. The above discussion clearly indicates that a single application of φ produces non-integrable spikes if $\gamma > 1$.

Chapter 5

Generalized space-splitting algorithm for hyperbolic systems with multidimensional unstable manifolds

Accurate approximations of the change of a system's output and its statistics with respect to the input are highly desired in computational dynamics. Ruelle's linear response theory provides breakthrough mathematical machinery for computing the linear response of chaotic dynamical systems. In this chapter, we propose an algorithm for sensitivity analysis of discrete chaos with an arbitrary number of positive Lyapunov exponents. We combine the concept of perturbation space-splitting, which regularizes Ruelle's original expression, together with measure-based parameterization of the expanding subspace. We use these tools to rigorously derive trajectory-following recursive relations that converge exponentially fast, and construct a memory-efficient Monte Carlo scheme for derivatives of the output statistics. Thanks to the regularization and lack of simplifying assumptions on the system's behavior, our method is immune to the common problems of other popular methods such as the exploding tangent solutions and unphysical shadowing directions. We provide a ready-to-use algorithm, analyze its complexity, and demonstrate several numerical examples of sensitivity computation using physically-inspired low-dimensional systems.

We acknowledge that this chapter was first published in *SIAM Journal*

5.1 Preliminaries

The main purpose of this work is to generalize the space-split algorithm, which was originally derived for systems with one-dimensional unstable manifolds. In other words, we seek a direct numerical method based on the linear response theory applicable to chaotic systems with an arbitrary number of positive LEs. The natural appeal of S3 is that it averages recursively-generated data along a trajectory and rigorously converges to the true solution in time. In the context of approximating the linear response, a "direct method" means that the only significant source of error is associated with Monte Carlo averaging. Thus, the crux of the space-splitting approach are recursive and converging relations for different quantities that arose from both the perturbation vector splitting and partial integration. The SRB density gradient, which represents an unstable derivative of the SRB measure, is one such a quantity. Indeed, as argued in Chapter 3, Lebesgue integration by parts requires knowledge of the measure derivative itself [160]. In the case of one-dimensional chaos, this iterative relation directly follows from the measure preservation property involving the Frobenius-Perron operator [158], which we showed in the appendices of Chapter 2 and Chapter 4. If the unstable manifold is geometrically more complex, it is convenient to apply the *measure-based* parameterization of the unstable subspace, directly relating directional derivatives of the SRB measure with the coordinate chart using quantile functions (inverse CDFs). In the case of the unit-speed parameterization used in [31], such a straightforward relation is not possible, which dramatically complicates the generalization of the iterations derived in that work to systems with more than one LE. This concept was demonstrated in the context of low-dimensional manifolds described by simple differential equations in [160], and later used to construct an ergodic-averaging scheme for the SRB density gradient [162]. The manifold descrip-

tion based on quantile functions is invariant to linear coordinate transformations and facilitates the iterative computation of directional derivatives of the SRB measure and other differentiable quantities by solving a collection of regularized tangent equations [162]. Here, we combine these two major concepts, i.e., perturbation space-splitting and measure-based parameterization, to derive a general Monte Carlo scheme for parametric derivatives of long-time averages defined by Eq. 1.1, in chaotic systems of arbitrary dimension. Instead of constructing expensive covariant bases of unstable manifolds to represent the perturbation splitting (see Chapter 5 of [28]), we compute cheaper orthogonal bases, which were also used in [127] to orthogonally regularize one of the linear response contributions. They align well with our algorithm and simplify the derivation of all critical quantities. The path we take in this chapter significantly differs from that of [127], since here we rely on a different representation of perturbation splitting, manifold description, and compute the SRB density gradient directly.

The major contributions of this chapter are the following:

- Applying the measure-based parameterization of unstable manifolds proposed in Chapter 3 [162] and the particular form of space-splitting from [31], we generalize the perturbation splitting using orthogonal bases and systematically derive a collection of trajectory-following recursions for unstable directional derivatives of the splitting coefficients and regularized tangent solutions.

- Based on the algebraic structure of the coordinate transformation matrices, its derivatives and orthogonality of Lyapunov bases, we analytically infer the key relation between the derivatives of Lyapunov vectors and second derivatives of the coordinate chart. The method for differentiating unstable bases is a by-product of our main algorithm and is applicable to many other problems.

- We analytically demonstrate the exponential convergence of the generalized trajectory-following recursions for derivatives of regularized tangent equations (derived in this chapter) and acceleration vectors (derived in [162]). We argue that the convergence of all other quantities follows from these two.

- We present several numerical examples that include the convergence analysis and comparison of S3 approximations of sensitivities against finite differences.

5.2 General Ruelle's linear response formula: fundamental aspects, practical consequences and S3

The purpose of this section is to review the linear response formula for discrete chaotic systems derived by Ruelle in [148, 150], introduce basic concepts from the dynamical systems theory that are critical in this work, and explain the notation convention used throughout this chapter. In addition, we present the concept of space-splitting, generalize it to systems of arbitrary dimension, and derive the stable contribution of the linear response.

Let us consider a parameterized discrete system,

$$x_{k+1} = \varphi(x_k; s) := \varphi(x_k), \tag{5.1}$$

governed by a diffeomorphism $\varphi : M \to M$, $M \subset \mathbb{R}^n$, $n \in \mathbb{Z}^+$, $s \in \mathbb{R}$, $k \in \mathbb{Z}$. This chapter focuses on parameterized systems and therefore all derivatives of φ also depend on s. To simplify the notation, the dependence on s shall not be explicitly indicated.

Notation conventions. This chapter intensively uses subscripts, superscripts, and a combination of both, that might have different meanings. Before we continue the technical discussion, we first explain the notation conventions that shall consistently be used in the remainder of this work:

- Subscripts involving $k + i$, $i \in \mathbb{Z}$, indicate that the underlying quantity is evaluated respectively at $(k+i)$-th time step of a random trajectory. It applies to all quantities appearing in this chapter, including map Jacobians and Hessians. For example, the observable J evaluated at k-th time step is denoted by $J(x_k) = J_k$. Analogously, the Jacobian of φ evaluated at x_k is denoted by $D\varphi_k$. Non-recursive equations without the "k" subscript hold at any point on the attractor.

Subscripts of this type should never appear inside integrals over M, because all possible points in M are included while integrating.

- Subscripts involving a number or symbol different than the one above have their own unique meaning. All of them are explained immediately after the first occurrence.

- Superscripts without round brackets indicate a particular representative of a family of scalars or arrays. For example, basis vectors of a certain m-dimensional subspace will be denoted by q^1, q^2, ..., q^m. They should not be confused with exponentiation. In this chapter, exponentiation is applied to the two scalars: m and n.

- Superscripts with round brackets refer to a particular component of an array. There might be more than one symbol in the superscript, which implies the array involves more dimensions. The colon symbol (:) is occasionally used in the superscript and means "all" components. For example, the j-th component of the i-th basis vector will be denoted by $(q^i)^{(j)}$. The i-th column of a matrix A is simply $A^{(:i)}$, while $A^{(ij)}$ represents the entry from the i-th row and j-th column of that matrix.

- Standard T and -1 symbols are used in superscripts to denote the transposition and inversion of a matrix, respectively.

Since we consider chaotic maps, φ has at least one positive Lyapunov exponent. As mentioned in Section 5.1, chaoticity is manifested by the exponential growth of the tangent solutions, which represent infinitesimal perturbations of the primal solution. The rate of growth of the tangent solutions is determined by LE values. We assume System 5.1 is *ergodic*, which implies the long-time average of a smooth observable $J : M \to \mathbb{R}$ can be computed in two distinct ways: 1) through averaging of the time series $\{J_0, J_1, J_2, ...\}$ generated along a trajectory, or 2), as an expected value of J with respect to the SRB measure μ [195]. Moreover, in ergodic systems, the long-time statistics do not depend on the choice of the μ-typical initial condition x_0. The SRB

measure is an invariant and physical measure that describes the statistical behavior of the system. This quantity is known to be absolutely continuous on the unstable manifold only. In the stable direction, across the expanding subspace, μ is generally singular with respect to the Lebesgue measure [195, 162]. The basic assumption of Ruelle's theory is uniform hyperbolicity. This property means that the tangent space decomposes into stable and unstable covariant subspaces everywhere on the attractor. The dimension of the latter is the same as the number of positive Lyapunov exponents and is fixed along any random trajectory [92]. Members of these two subspaces are solutions to the homogeneous tangent equations and their norms uniformly decay and grow in time at an exponential rate, respectively. Recall that there also exists a neutral subspace aligned with the flow in continuous-time systems (ODEs), which are beyond the scope of this work.

Under the assumption of uniform hyperbolicity, Ruelle rigorously showed that

$$\frac{d\langle J\rangle}{ds} = \sum_{t=0}^{\infty} \int_M D(J \circ \varphi_t) \cdot \chi \, d\mu, \tag{5.2}$$

where D represents the differentiation (gradient) operator in phase space, while $\chi = \partial_s \varphi \circ \varphi_{-1}$ is the map perturbation vector and $\varphi_t = \varphi(\varphi_{t-1})$, $\varphi_0(x) = x$, $\partial_s(\cdot) := \partial(\cdot)/\partial_s$. This result assumes that the observable J does not depend on the parameter s. If it were otherwise, the expected value of $\partial_s J$ would need to be added on the right-hand side (RHS). Note that Eq. 5.2 is identical to Eq. 4.16, which suggests that the original linear response expression is agnostic to the number of LEs. Although Eq. 5.2 is rigorously true for uniformly hyperbolic diffeomorphisms, a modified version of Ruelle's expression has been empirically shown to be valid in statistical mechanics systems that violate this basic assumption [113]. Indeed, various high-dimensional systems describing complex physical phenomena behave as if they were uniformly hyperbolic [59, 153]. Therefore, Ruelle's expression could potentially be applied to various real-world chaotic systems encountered in science and engineering.

Although Eq. 5.2 provides a closed-form expression for the desired quantity, its direct evaluation is impractical. One could naïvely approximate each integral of

the infinite series using the Ergodic Theorem, i.e., by taking the long-time average of the integrand along a trajectory. Note, however, that the integrand involves a derivative of the observable J evaluated t time steps ahead in time. Differentiating this composite function requires computing the product of the map Jacobians $D\varphi$ evaluated at t consecutive time steps. Owing to the butterfly effect, this product grows exponentially with t at the rate proportional to the largest LE, λ_1. Assuming J is sufficiently smooth, one can rigorously show that Ruelle's formula is equivalent to

$$\frac{d\langle J\rangle}{ds} = \int_M DJ \cdot u \, d\mu, \tag{5.3}$$

where $u = \partial_s x$, and

$$u_{k+1} = D\varphi_k \, u_k + \chi_{k+1}, \quad u_0 = 0. \tag{5.4}$$

According to our notational convention,

$$\chi_{k+1} = \partial_s \varphi \circ (\varphi_{-1}(x_{k+1})) = \partial_s \varphi(x_k).$$

Note that Eq. 5.4 is obtained by differentiating Eq. 5.1 with respect to s. The computation of the sequence $\{u_0, u_1, ...\}$ through the above inhomogeneous tangent equation (Eq. 5.4) is ill-conditioned, because

$$u_k = \sum_{i=1}^{k-1} \left(\prod_{j=i}^{k-1} D\varphi(x_j) \right) \chi(x_i). \tag{5.5}$$

Eq. 5.5 implies that

$$\|u_k\| \sim \mathcal{O}(\exp(\lambda_1 k)).$$

The equivalence of Eq. 5.3-5.4 and Eq. 5.2 directly follows from the chain rule applied to an ergodic system. The problem of exploding tangent solutions is absent only in the two following scenarios: 1) $\lambda_1 < 0$ (non-chaotic system), or 2), the tangent solution u is orthogonal to the unstable manifold. The first scenario is trivial, because one can use Eq. 5.4 without further modifications to estimate sensitivities, while the second scenario is highly unlikely.

A natural remedy for the exploding product of Jacobians is the integration by parts applied to the original Ruelle's expression. In the case of Lebesgue integrals, however, one also needs to differentiate the measure itself. Note that χ does not generally belong to the unstable subspace and, therefore, the direct partial integration is prohibited. To circumvent this mathematical difficulty, we apply the main idea of the space-split sensitivity (S3) method [31], which was motivated by Ruelle's original work [148, 150], to decompose the perturbation vector χ into two terms, $\chi = \chi_1 + \chi_2$, which we call the unstable and stable contribution, respectively. The original S3 splitting was constructed such that χ_1 is tangent to the one-dimensional expansive subspace enabling partial integration. To guarantee the differentiability of both terms on unstable manifolds, χ_1 must be expressed by a linear combination of basis vectors of the expanding subspace. Despite its name, the other term, $\chi_2 = \chi - \chi_1$, does not strictly belong to the stable subbundle.

Motivated by this concept, we require that χ_1 belongs to the m-dimensional unstable manifold at any typical point on the attractor. Recall that here we assume φ has $m \geq 1$ positive LEs. Let $q^i(x_k) := q^i_k$, $i = 1, ..., m$ denote an orthonormal basis of the unstable manifold at $x_k \in M$. Thus, we propose the following perturbation splitting,

$$\chi = \chi_1 + \chi_2 = \left(\sum_{i=1}^{m} c^i \, q^i \right) + \left(\chi - \sum_{i=1}^{m} c^i \, q^i \right), \tag{5.6}$$

where c^i is a scalar differentiable on unstable manifolds. This decomposition is local, i.e., both the unstable basis and scalar coefficients vary from point to point on the attractor. If we plug Eq. 5.6 back to Ruelle's formula, we observe that the first term can now be integrated by parts, regardless of the choice of c^i. Indeed, the unstable contribution involves m directional derivatives of $J \circ \varphi_t$ in the directions indicated by the selected orthonormal basis of the unstable manifold. The second term, i.e., the stable contribution, can be rigorously re-expressed as a single Lebesgue integral, analogously to Eq. 5.3,

$$\sum_{t=0}^{\infty} \int_M D(J \circ \varphi_t) \cdot \chi_2 \, d\mu = \int_M DJ \cdot v \, d\mu, \tag{5.7}$$

where v is a solution to the following *regularized* tangent equation,

$$v_{k+1} = D\varphi_k\, v_k + \chi_{k+1} - \sum_{i=1}^{m} c^i_{k+1}\, q^i_{k+1}, \quad v_0 = 0, \tag{5.8}$$

which is derived in the same fashion as its original counterpart in Eq. 5.4. The proposed splitting gives us freedom in choosing the values of $c^i,\ i = 1, ..., m$. By imposing the following set of m scalar constraints,

$$v \cdot q^i = 0,\ i = 1, ..., m, \tag{5.9}$$

we fix all c^i and, simultaneously, enforce the tangent solution v to be orthogonal to the unstable manifold everywhere along a trajectory. This guarantees that the norm of v does not increase exponentially in time. Therefore, by combining Eq. 5.8-5.9, we obtain a linear system with $n + m$ equations and the same number of unknowns (n components of v and m scalars c^i). Using the fact $q^i \cdot q^j = 1$ if $i = j$ and $q^i \cdot q^j = 0$ otherwise, the orthogonality constraint (Eq. 5.9) can be enforced by setting

$$c^i_{k+1} = q^i_{k+1} \cdot \left(D\varphi_k\, v_k + \chi_{k+1}\right),\ i = 1, ..., m. \tag{5.10}$$

We refer the reader to [31] and Chapter 6 to see analogous, yet much more complicated, expressions involving splitting coefficients if a covariant basis is used and the orthogonality constraint is violated. Due to the ergodicity and uniform hyperbolicity of φ, the RHS of Eq. 5.7 can be approximated as a finite-time average of $DJ \cdot v$. Assuming the integrand is Hölder continuous, the ergodic-averaging error is bounded above as follows,

$$\left| \int_M DJ \cdot v\, d\mu - \frac{1}{N} \sum_{k=0}^{N-1} DJ_k \cdot v_k \right| \leq \frac{C_1}{N} + C_2 \sqrt{\frac{\log \log N}{N}}. \tag{5.11}$$

for some real constants $C_1 > 0$ and $C_2 > 0$, regardless of the value of m. The reader is referred to [31] (Proposition 9.1) for a proof of the preceding statement. If N is large, the second term of the above upper bound determines the actual convergence

rate as it decays slower. Note also that $\log\log N$ grows very slowly with N, which means that $\sqrt{(\log\log N)/N}$ and $\sqrt{1/N}$ are approximately parallel for large values of N. This further implies that the convergence rate of the "stable" integrand is very similar to the one of a typical Monte Carlo process.

Therefore, the computation of the stable contribution requires solving a constrained tangent equation (Eq. 5.8-5.9) and taking the time average of $DJ \cdot v$, which is recursively computed as we move along an orbit initiated at a μ-typical point. Note that the dimension of the unstable manifold has little impact on the complexity of that part of the algorithm. To enforce all the orthogonality constrains, one naturally needs to pre-compute the basis of the unstable manifold everywhere along the trajectory. We postpone the discussion on numerical procedures for approximating the unstable basis vectors until Section 5.3, as they are essential ingredients of the unstable contribution.

Key point. The perturbation vector space-splitting proposed in Eq. 5.6 is absolutely critical in regularizing Ruelle's linear response expression for the sensitivity of a chaotic system. Indeed, it guarantees the differentiability of both terms on the unstable subbundle, and allows to partially integrate the term involving the unstable component of χ and apply the conventional tangent equation approach together with recursive orthogonalization (stabilization) to approximate the stable term. The procedure for approximating the latter is largely agnostic to the dimension of the unstable manifold. This is not the case for the unstable contribution, which is the main focus of the following section.

5.3 Computation of the unstable contribution

5.3.1 Reviewing critical concepts: measure-based parameterization, integration by parts, and SRB density gradient

To approximate the desired sensitivity, $d\langle J\rangle/ds$, one needs to sum up two contributions, the unstable and stable terms, as defined in the previous section (Eq. 5.6).

Since the former contains a component of the perturbation vector that is a member of the unstable manifold, we are allowed to apply partial integration to move the differentiation operator away from the troublesome composite function. However, we can do so only on unstable manifolds, because the SRB measure μ is generally non-smooth on M. Therefore, an extra step involving measure disintegration is required before applying the integration by parts. Let us consider a measurable partition U of M aligned with the geometry of unstable manifolds. Thus, for any Borel subset $B \subset M$,

$$\mu(B) = \int_{M/U} \tilde{\mu}_x(B \cap U_x)\, d\hat{\mu}(x), \tag{5.12}$$

where $\hat{\mu}$ denotes the quotient measure defined by the partition U, while $\tilde{\mu}_x$ represents the SRB measure conditioned on the unstable manifold U_x that contains $x \in M$. Intuitively, Eq. 5.12 means that the measure of B can be computed by summing conditional measures of local intersections weighted by the likelihood of each partition member. Using Eq. 5.12, the unstable contribution can be expressed as follows,

$$\sum_{t=0}^{\infty} \int_M D(J \circ \varphi_t) \cdot \chi_1 \, d\mu = \sum_{t=0}^{\infty} \sum_{i=1}^{m} \int_M c^i \, \partial_{q^i}(J \circ \varphi_t)\, d\mu, \tag{5.13}$$

where $\partial_{q^i}(\cdot) := D(\cdot) \cdot q^i$ is a short-hand notation for the directional derivative in the direction of the i-th basis vector. Within each Lebesgue integral of the above double sum, we apply measure disintegration and then integrate by parts on U_x [152],

$$\begin{aligned} &\int_M c^i \, \partial_{q^i}(J \circ \varphi_t)\, d\mu = \\ &\int_{M/U} \int_{U_x} c^i \, \partial_{q^i}(J \circ \varphi_t)\, d\tilde{\mu}_x \, d\hat{\mu}(x) = \int_{M/U} \int_{U_x} c^i \, \partial_{q^i}(J \circ \varphi_t)\, \tilde{\rho}_x \, d\tilde{\omega}_x \, d\hat{\mu}(x) = \\ &- \int_{M/U} \int_{U_x} J \circ \varphi_t \left(\partial_{q^i} c^i + c^i \frac{\partial_{q^i} \tilde{\rho}_x}{\tilde{\rho}_x} \right) d\tilde{\mu}_x \, d\hat{\mu}(x) + \text{B.T.}, \end{aligned} \tag{5.14}$$

where $\tilde{\rho}_x$ and ω_x respectively represent the density of the conditional measure and the natural volume form, both defined on U_x. The second term on the RHS of Eq. 5.14 represents the boundary term, denoted by B.T., which can be expressed as the divergence of a smooth field on unstable manifolds. In uniformly hyperbolic systems,

this term rigorously vanishes according to Theorem 3.1(b) of [148]. The reader is also referred to Section 3.6 for a more intuitive explanation of this counter-intuitive cancellation. We eventually obtain a new integral that involves two quantities, $b^{i,i}$ and g^i, defined as follows,

$$b^{i,j} := \partial_{q^j} c^i, \; g^i := \frac{\partial_{q^i} \tilde{\rho}}{\tilde{\rho}} = \partial_{q^i} \log \tilde{\rho}. \tag{5.15}$$

The computation of these two quantities is the actual price for the regularization of the original Lebesgue integrals. The latter is known in the literature as the SRB density gradient has thoroughly been described in Chapter 3. Recall that it reflects a relative measure change along an unstable manifold and thus its value is independent from its corresponding quotient measure. An efficient trajectory-driven procedure for the computation of g has been summarized by Algorithm 1, which we will directly use in this chapter. Therefore, the final missing ingredients are the directional derivatives of c^i. While Eq. 5.14 involves $b^{i,j}$ with $i = j$ only, we will show that all possible m^2 representatives of $b^{i,j}$ are required in the algorithm.

We intend to utilize the measure-based parameterization of unstable manifolds proposed in Chapter 3. For the sake of clarity, we review the major points here. Let us consider a family of smooth charts $x_k(\xi) : [0,1]^m \to U_k$, where U_k is a partition member that is crossed by the trajectory at time step k. Consequently, U_k is an m-dimensional unstable manifold such that $x_k \in U_k$. The measure-based parameterization is defined such that the SRB measure of any Borel subset $V \in [0,1]^m$ satisfying $x_k(V) = B_k \subset U_k$ is related to the corresponding SRB density through

$$\tilde{\mu}_k(V) = \int_{B_k} \tilde{\rho}_k \, d\omega_k. \tag{5.16}$$

Notice that we replaced x with k in the subscript of the local/conditional quantities to relate them with a particular point on a random trajectory. Indeed, for a given trajectory initiated at a randomly chosen μ-typical initial condition, k uniquely determines a point on the attractor. The major benefit of this type of description is a straightforward relation between the parametric gradient of $x(\xi)$, denoted by $\nabla_\xi x$,

and the conditional SRB density $\tilde{\rho}$. In particular,

$$\tilde{\rho}(x(\xi)) \,|\det R(x(\xi))| = 1 \tag{5.17}$$

for any $\xi \in [0,1]^m$, where R is an $m \times m$ invertible matrix obtained through the QR factorization (orthonormalization) of the chart gradient,

$$\nabla_\xi x(\xi) = Q(x(\xi))\, R(x(\xi)), \tag{5.18}$$

where $Q^T Q = I$ and R is an upper-triangular matrix containing projections of the columns of $\nabla_\xi x$ onto its orthonormal basis stored in the Q matrix. Eq. 5.17 is a general representation of the measure conservation in a nonlinear transformation from a uniform (constant) to non-uniform distribution.

The crux of the SRB density gradient computation, as explained in [160], relies on a recursive computation of the first- and second-order derivatives of the coordinate chart. These recursive formulas are derived by taking derivatives of the original system (Eq. 5.1) and applying the chain rule on smooth manifolds. However, note that the naïve computation of the chart gradient is ill-conditioned, because $\nabla_{\xi_0} x_k = (\prod_{i=0}^{k-1} D\varphi_i)\nabla_{\xi_0} x_0$ grows in norm exponentially fast as discussed in Section 5.2, where ξ_0 represents the initial, i.e., the one chosen at $k = 0$, parametric coordinate system. In their recent work [162], the authors proposed a step-by-step orthonormalization of the chart gradient through a recursive update of the coordinate system using the following linear transformation,

$$\xi_{k+1} = R(x_{k+1}(0))\, \xi_k. \tag{5.19}$$

Applying this coordinate change in a step-by-step manner, we ensure the parametric gradient computed with respect to the new coordinates is orthogonal at the origin, i.e., at $\xi = 0$. In practice, this requires performing the QR factorization every time step, where Q contains the orthogonal basis, while R is used to transform coordinates. In the orthogonalized coordinate system at $\xi = 0$, the SRB measure gradient is can

be expressed as a projection of the unstable basis onto a set of the acceleration (i.e., second derivatives of the chart) vectors denoted by a [162],

$$g^i = -q^j \cdot \partial_{\xi^{(i)}} \partial_{\xi^{(j)}} x := -q^j \cdot a^{i,j}, \tag{5.20}$$

where the repeated indices imply summation per Einstein's convention, while a satisfies the following recursion,

$$a^{i,j}_{k+1} = \left(D^2 \varphi_k(q^p_k, q^q_k) + D\varphi_k \, a^{p,q}_k \right) (R^{-1}_{k+1})^{(pi)} \, (R^{-1}_{k+1})^{(qj)}. \tag{5.21}$$

The product $D^2\varphi(a,b)$ represents the contraction of the Hessian of φ against two vectors, a and b. This operation outputs a vector whose i-th component equals $(D^2\varphi(a,b))^{(i)} = \partial_{x^{(p)}} \partial_{x^{(q)}} \varphi^{(i)} \, a^{(p)} \, b^{(q)}$. The choice of $\xi = 0$ does not restrict our algorithm to a certain trajectory. Note that one can freely stretch or shrink the feasible space of ξ such that the preimage of the initial state x_0 is $\xi_0 = 0$. From now on, all equations will be evaluated at $\xi = 0$ and, in order to further simplify the notation, the argument of coordinate charts will be skipped.

5.3.2 Derivation of recursions for the missing directional derivatives and convergence study

We shall now analyze the convergence of the recursive algorithm for the SRB density gradient. Based on the above description, the iterative computation of the basis matrix Q involves two steps, i.e., left-multiplying Q by the Jacobian matrix followed by QR factorization of the obtained matrix product. This implies that the basis matrix at the k-th time step equals

$$Q_k = D\varphi_{k-1} \ldots D\varphi_0 \, Q_0 \, R_1^{-1} \ldots R_k^{-1}. \tag{5.22}$$

If one replaces Q_0 with any arbitrary matrix that is bounded in norm, then the process described by Eq. 5.22 is guaranteed to converge at an exponential rate if φ is a

uniformly hyperbolic diffeomorphism [90]. It means that in ideally chaotic systems one can generate the basis vectors of unstable manifolds, a.k.a. the backward Lyapunov vectors, by running the trajectory-driven iteration described above. In such systems, the product of the inverses of subsequent R matrices decays in norm at an exponential rate or faster as k increases [51], i.e.,

$$\|R_1^{-1}\, R_2^{-1} ... R_k^{-1}\| \leq \exp(-ck),\ c > 0. \tag{5.23}$$

These matrices counterbalance the exploding product of Jacobians along typical trajectories.

Given this remarkable behavior, we conclude that the iterative process for a must also converge. To see that, let us consider a difference between two approximations of a along a single trajectory assuming the basis vectors are the same in both the iterations, labelled as 1 and 2,

$$a_{k+1,1}^{i,j} - a_{k+1,2}^{i,j} := \delta a_{k+1}^{i,j} = D\varphi_k\, \delta a_k^{p,q}\, (R_{k+1}^{-1})^{(pi)}\, (R_{k+1}^{-1})^{(qj)}. \tag{5.24}$$

We observe Eq. 5.24 describes the evolution of the differences of acceleration vectors along a trajectory. This equation implies that the differences are recursively left-multiplied by the map Jacobian combined with a double contraction against the R matrix. Note that the RHS of Eq. 5.24 can be viewed as a two-step algebraic process. In the first step, one computes n matrix products

$$(R_{k+1}^T)^{-1}\, \delta A_k^i\, (R_{k+1})^{-1},\ i = 1, ..., n,$$

where δA_k^i is an $m \times m$ matrix that contains the i-th components of $\delta a_k^{p,q}$, $p, q = 1, ..., m$. Subsequently, the resulting n-dimensional vectors (there are m^2 of them) are left-multiplied by the same Jacobian matrix.

It is critical to note that the two algebraic steps are commutative, which means that we are allowed to take the initial differences $\delta a_0^{p,q}$, $p, q = 1, ..., m$, left-multiply them by a product of k Jacobians and then recursively compute the double contrac-

tions against k inverses of R. Note also that the double contraction can be split into two single ones, which further implies that we could, for example, recursively left-multiply the difference vector by a Jacobian with a single contraction, which is equivalent to replacing $(R_{k+1}^{-1})^{pi}$ with a Kronecker delta δ^{pi}. To compute the true solution at $(k+1)$-th time step, the obtained vectors still need to be recursively contracted against $k+1$ inverses of R once more. The purpose of this discussion is to argue that if we replace one contraction with an identity operation in Eq. 5.24, we effectively obtain a recursion equivalent to the one in Eq. 5.22, which produces vectors with norms of the order $\mathcal{O}(1)$. The second contraction appearing in the original version of Eq. 5.22 means that these vectors are left-multiplied by the product $R_{k+1}^{-1}...R_1^{-1}$ whose induced norm uniformly approaches 0 at an exponential rate.

Therefore, if the iteration defined by Eq. 5.22 exponentially converges to the true solution regardless of the choice of Q_0, which is true for uniformly hyperbolic systems, the recursion for a (Eq. 5.21) also converges to its true value at an exponential rate. The ultimate implication is that the iterative algorithm for the SRB density gradient g does not depend on the initial guess and its true value can be obtained after a moderately small number of iterations.

The remaining part of this section focuses on recursive computation of b, which is the final term required to evaluate the RHS of the regularized unstable contribution in Eq. 5.14. Recall that b equals a parametric derivative of the scalars appearing in the constrained tangent equation. These scalars are directly computed using Eq. 5.10. Recall also that we describe the unstable manifold using a smooth chart $x(\xi)$ with a linearly re-scaled coordinate system that ensures the orthogonality of its gradient at $\xi = 0$ as introduced above. Thus, by differentiating Eq. 5.10, we obtain an explicit formula for b,

$$\begin{aligned} b_{k+1}^{i,j} = \partial_{q_{k+1}^j} c_{k+1}^i \overset{\|\partial_{\xi_{k+1}^{(j)}} x_{k+1}(0)\|=1}{=\!=\!=\!=\!=\!=\!=} \partial_{\xi_{k+1}^{(j)}} c_{k+1}^i = \\ \partial_{\xi_{k+1}^{(j)}} q_{k+1}^i \cdot f_k + q_{k+1}^i \cdot \partial_{\xi_{k+1}^{(j)}} f_k := p_{k+1}^{i,j} \cdot f_k + q_{k+1}^i \cdot \partial_{\xi_{k+1}^{(j)}} f_k, \end{aligned} \tag{5.25}$$

where

$$f_k := D\varphi_k\, v_k + \chi_{k+1} = D\varphi_k\, v_k + \partial_s \varphi_k.$$

While the recipe for f and q has already been discussed, we still require two more quantities, p and parametric derivative of f, in order to complete the algorithm. We first focus on p, which represents the directional derivative of a backward Lyapunov vector at the origin of the updated coordinate system. In general, however, $p^{i,j}$ does not equal $a^{i,j}$. The latter is defined as the second parametric derivative of the chart evaluated at the origin. Our new quantity p, on the other hand, is defined as the parametric derivative of q also evaluated at the origin. Consider a 1D smooth function $h(e) : \mathbb{R} \to \mathbb{R}$ and coordinate transformation $y(e)$. The difference between a and p is equivalent to the difference between $h''(y(e))$ and $(h'(y(e)))'$ at a particular instance of e. Therefore, to relate a an p, one also needs to differentiate the re-scaling factor represented by the R matrix. This relationship can be found by differentiating Eq. 5.18 with respect to the i-th chart coordinate,

$$\partial_{\xi^{(i)}}(\nabla_\xi x) = (\partial_{\xi^{(i)}} Q)\, R + Q\, (\partial_{\xi^{(i)}} R), \tag{5.26}$$

which implies that

$$(\partial_{\xi^{(i)}} R)\, R^{-1} = Q^T\, \partial_{\xi^{(i)}}(\nabla_\xi x)\, R^{-1} - Q^T\, (\partial_{\xi^{(i)}} Q), \tag{5.27}$$

for any $\xi \in [0,1]^m$. At the origin of the orthonormalized coordinate system, however, the R matrix equals the identity by construction and thus

$$\partial_{\xi^{(i)}} R = Q^T\, A^i - Q^T\, P^i, \tag{5.28}$$

where A^i and P^i respectively contain second parametric derivatives of the chart and first parametric derivatives of basis vectors, both evaluated at the origin. Note that $(A^i)^{(:j)} := a^{j,i} = a^{i,j}$ assuming $x(\xi)$ is sufficiently smooth and, analogously, $(P^i)^{(:j)} := p^{j,i}$. Although Eq. 5.28 provides an explicit relation between A and P, we still need

more information to compute the latter as the parametric derivative of R is unknown.

The missing puzzle piece is hidden is the structure of the matrices appearing in Eq. 5.27. Indeed, the LHS of that equation is always upper-triangular since it is a product of two other upper-triangular matrices, while the second term on the RHS must be skew-symmetric, because

$$\partial_{\xi^{(i)}} \left(Q^T \, Q \right) = \partial_{\xi^{(i)}} I = 0 \implies \left(Q^T \, \partial_{\xi^{(i)}} Q \right)^T + \left(Q^T \, \partial_{\xi^{(i)}} Q \right) = 0.$$

Therefore, we conclude that

$$(\partial_{\xi^l} R)^{(ij)} = \begin{cases} q^i \cdot a^{j,l} & \text{if } i = j, \\ q^i \cdot a^{j,l} + q^j \cdot a^{i,l} & \text{if } i < j \\ 0 & \text{otherwise.} \end{cases} \tag{5.29}$$

We now combine Eq. 5.29 and Eq. 5.26 to infer an explicit expression for derivatives of backward Lypaunov vectors at the origin,

$$p^{i,j} = a^{i,j} - q^l (\partial_{\xi^j} R)^{(li)}. \tag{5.30}$$

Note that the computation of p requires only the knowledge of a and Q, both of which are integral components of the algorithm for the SRB density gradient g [162]. Therefore, the procedure for g extended by the two above equations, Eq. 5.29 and Eq. 5.30, enables recursive computation of p along a typical trajectory. If the procedure for g converges exponentially fast as argued above, the same is true of its extended version. Eq. 5.29 clearly indicates that, in general, $p^{i,j} \neq p^{j,i}$ if $i \neq j$. That lack of symmetry requires us to compute all m^2 different p vectors to advance the full algorithm in time, which will be clear at the end of this section. Note that Eq. 5.30 reduces to $p = a + g\, q$ if $m = 1$, where p, a, q are vectors, while g is a scalar. Note also that $p = a$ only if the measure is distributed uniformly. This simplified equation was also identified in [31] (between Eq. 6.25 and 6.26), but nothing insightful could be inferred from that simple 1D case. We implemented this simplification in the attached

code for the baker's map (see Section 5.5). We also acknowledge that Eq. 5.30 is generally important in other applications spanning the fields of differential geometry and dynamical systems (see [33] and references therein).

The final task in the derivation of the full algorithm is to apply the chain rule in the second term of the RHS of Eq. 5.25. Notice that at $\xi_{k+1} = \xi_k = 0$, one can directly change variables of the differentiation because $R_{k+1}^{-1} = \partial\xi_k/\partial\xi_{k+1}$, which implies that

$$\nabla_{\xi_{k+1}} f_k = \nabla_{\xi_k} f_k \, R_{k+1}^{-1}, \tag{5.31}$$

where the i-th column of $\nabla_{\xi_k} f_k$ can be expanded as follows,

$$\partial_{\xi_k^{(i)}} f_k = D^2\varphi_k(v_k, q_k^i) + D\varphi_k \, w_k^i + D\partial_s\varphi_k \, q_k^i. \tag{5.32}$$

The matrix $D\partial_s\varphi_k$ represents the Jacobian of the map differentiated with respect to the scalar s and evaluated at time k. The new quantity, w_k^i, represents a parametric derivative of the regularized tangent vector, i.e., $w_k^i := \partial_{\xi_k^{(i)}} v_k$, and is recursively computed in the following way,

$$w_{k+1}^i = (\nabla_{\xi_{k+1}} f_k)^{(:i)} - b_{k+1}^{l,i} \, q_{k+1}^l + c_{k+1}^l \, p_{k+1}^{l,i}. \tag{5.33}$$

This formula is obtained through parametric differentiation of Eq. 5.8. We can now observe that the entire set of m^2 scalars $b^{i,j}$ and m^2 vectors $p^{i,j}$ are necessary in order to advance the iteration for w^i. While b appears in the recursion for w and vice versa, there is no need to construct large linear systems to find both the quantities. Indeed, Eq. 5.25, 5.31, 5.32 indicate that in order to find b at time $k+1$, we need all vectors w at the previous time k. Therefore, in our algorithm, we can sequentially compute all vectors/scalars in the following order: a, p, b and w, at every point along a trajectory.

We already discussed the convergence of the iterations for Q, a, and p. Our final task is the convergence analysis of the recursion for w (Eq. 5.33). Let $\delta W_k := (W_1)_k - (W_2)_k$ be the difference of two matrices containing all vectors w in their columns such that $(W_1)_k^{(:i)} := (w_1)_k^i$ and $(W_2)_k^{(:i)} := (w_2)_k^i$, while the labels 1 and 2

represent two different and randomly chosen initial conditions for the recursion of w. Therefore, $\|\delta W_0\| \neq 0$ in general. Using this notation and combining Eq. 5.33 and Eq. 5.25, we derive the following iteration for the difference matrix,

$$\delta W_{k+1} = \left(I - Q_{k+1}\, Q_{k+1}^T\right) D\varphi_k\, \delta W_k\, R_{k+1}^{-1}. \tag{5.34}$$

Note that the difference matrix is left-multiplied by another matrix that is orthogonal to the unstable manifold, because

$$Q^T \left(I - Q\, Q^T\right)(\cdot) = 0.$$

Therefore, the recursive application of the left-hand side operator to any initial difference matrix converges to zero in norm at an exponential rate, i.e.,

$$\|P_k\, P_{k-1} ... P_1\, \delta W_0\| \leq \exp(-ck), \quad P_k := \left(I - Q_{k+1}\, Q_{k+1}^T\right) D\varphi_k, \tag{5.35}$$

for some constant $c > 0$. To see this, let us consider a generic member of the tangent space, v, that is bounded in norm. Any such v can be linearly decomposed such that $v = C_u v + C_s v$ with $C_u v$ and $C_s v$ being its distinct components strictly belonging to the unstable and stable manifolds, respectively. Thus, $(I - Q\, Q^T) C_u\, v = 0$ and $(I - Q\, Q^T) C_s\, v = C_s\, v$. In addition, uniform hyperbolicity guarantees that the product $(\prod_{k=0}^{N} D\varphi_k) C_s v$ strictly belongs to the stable subspace (covariance property), while its norm is bounded above by $C\lambda^N \|v\|$ with $C > 0$ and $\lambda \in (0, 1)$ (uniform decay property). Note also that the initial difference δW_0 is bombarded by the product of the inverses of R matrices, which also decays with k exponentially fast per our discussion above.

Key point. The ultimate conclusion of this analysis is that all the recursions presented and derived in this section do not depend on initial conditions and their respective solutions converge to their true values exponentially fast. This discussion extends the convergence analysis presented in [31] to systems with multidimensional unstable manifolds.

Having the collection of converging iterative expressions for different quantities arising in the regularized version of the unstable contribution, the final step is to take the time average of the series generated along a typical trajectory. In particular, assuming the system is ergodic and combining Eq. 5.13-5.14, we approximate the unstable contribution through the following triple sum,

$$\sum_{t=0}^{\infty}\int_M D(J\circ\varphi_t)\cdot\chi_1\,d\mu \approx \frac{1}{N}\sum_{t=0}^{K}\sum_{k=0}^{N}\sum_{i=1}^{m} J_{k+t}\,\left(b_k^{i,i}+c_k^i\,g_k^i\right) := \sum_{t=0}^{K} U(N,t), \quad (5.36)$$

where T, K are some positive integers. Assuming all the quantities appearing in the above expression are Hölder continuous, the law of iterated logarithm applies and the truncated series approximating ergodic averages converge as $\mathcal{O}(\sqrt{\log\log N}/\sqrt{N})$, regardless of the value of m. This means that for a finite K [31],

$$\left|\sum_{t=0}^{K}\int_M D(J\circ\varphi_t)\cdot\chi_1\,d\mu - U(N,t)\right| \leq C(K)\sqrt{\frac{\log\log N}{N}}, \quad (5.37)$$

where $C(K)$ is a positive scalar that depends on K. It means that for a fixed K, the ergodic-averaging procedure for computing the unstable term approximately behaves as typical Monte Carlo procedure, analogously to its stable counterpart. However, note that by truncating the original infinite series, we approach a biased value that might be significantly different than the true one if K is chosen improperly. Recall also that due to the exponential decay of correlations, which is guaranteed in uniformly hyperbolic systems, the sum over t convergences to the true solution exponentially fast if ergodic averages are exact. In other words, by setting $N=\infty$ and summing over k first, the triple sum in the RHS of Eq. 5.36 rigorously converges to the true solution exponentially fast as $K\to\infty$.

Key point. Two practical conclusions follow form this analysis. First, both N and K should be sufficiently large to guarantee that both the Monte Carlo error and series truncation bias are small. Second, for a fixed and finite N^*, increasing K might actually increase the approximation error at some point. Indeed, notice that there exists a K^* for which $\sum_{t=K^*}^{\infty}\int_M D(J\circ\varphi_t)\cdot\chi_1\,d\mu$ is within $\mathcal{O}(\epsilon)$ from zero with $\epsilon>0$

[31]. However, this does not guarantee $C(K^*)\sqrt{(\log\log N^*)/N^*}$ is $\mathcal{O}(\epsilon)$. In practice, we always truncate a random trajectory and, consequently, we do not let the terms corresponding to larger values of K become sufficiently small.

The above discussion concludes the generalized space-split algorithm. A thorough summary of the entire algorithm, analysis of its computational complexity, practical implementation guidelines, and demonstration of several numerical examples are presented in the following two sections.

5.4 Space-split algorithm for multi-dimensional hyperbolic system

We now synthesize all derivations and analysis presented in Sections 5.2–5.3, and construct an algorithm for sensitivity computation of chaotic dynamical systems with an arbitrary number of degrees of freedom n and positive Lyapunov exponents m. Algorithm 2 is a summary of the space-split procedure in the form of a pseudocode.

Every iteration of the main time loop starts from updating the sums of the stable and unstable integrands, s and u, respectively (Lines 3-6). We disregard the first T data points to ensure that all quantities contributing to the final average are close to their respective true values up to the machine precision. Given that all recursions exponentially converge, the value of T is low compared to N, i.e., $T \ll N$.

Lines 7-11 are taken from [162], as they reflect all the steps necessary to compute the SRB density gradient. Note that this code chunk involves advancing m tangent equations (Line 7), QR factorization (Line 8), inverting the R matrix (Line 9), advancing $m^2/2$ second-order tangent equations (Line 10), and double contraction of $m^2/2$ acceleration vectors against the R^{-1} matrix (Line 11). Indeed, the most expensive stage of this chunk is Line 10, which costs $\mathcal{O}(n^3\,m^2)$ due to the presence of the third-order tensor (Hessian of φ) contracted against two different vectors. This is because for each component of the new n-dimensional vector $\tilde{a}$, one must compute and sum up n^2 different scalar products. Note also that the re-scaling stage (Line 11) involves

Algorithm 2: Space-split sensitivity (S3) algorithm for discrete systems

Input : N, K, T, n, m, $s = 0$, $u = 0$

Output: $d\langle J\rangle/ds \approx (s+u)/N$

Randomly generate: x_0, v_0, Q_0, $a_0^{i,j}$, w_0^i for all $i, j = 1, ..., m$;

for $k = 0, ..., N-1$ **do** `// main time loop`

  if $k \geq T$ **then**

    $s = s + DJ_k \cdot v_k$;

    $u = u - J_k\,(u_k + u_{k-1} + ... + u_{k-K+1})$;

  end

  $S_k = D\varphi_k\, Q_k$;

  QR-factorize S_k: $Q_{k+1}\, R_{k+1} = S_k$;

  Find the inverse of R_{k+1};

  for $i = 1, ..., m,\ j = 1, ..., i$ **do** $\tilde{a}_{k+1}^{i,j} = D^2\varphi_k(q_k^i, q_k^j) + D\varphi_k\, a_k^{i,j}$;

  for $i = 1, ..., m,\ j = 1, ..., i$ **do** $a_{k+1}^{i,j} = \tilde{a}_{k+1}^{p,q}\,(R^{-1})_{k+1}^{(pi)}\,(R^{-1})_{k+1}^{(qj)}$;

  for $i = 1, ..., m$ **do**

    for $p, q = 1, ..., m$ **do**

$$(\partial_{\xi_{k+1}^i} R_{k+1})^{(pq)} = \begin{cases} q_{k+1}^p \cdot a_{k+1}^{p,i}, & \text{if } p = q \\ q_{k+1}^p \cdot a_{k+1}^{q,i} + q_{k+1}^q \cdot a_{k+1}^{p,i}, & \text{if } p < q \\ 0, & \text{otherwise} \end{cases} ;$$

    end

    $g_{k+1}^i = -\mathrm{tr}(\partial_{\xi_{k+1}^i} R_{k+1})$;

  end

  $f_k = D\varphi_k\, v_k + \chi_{k+1}$;

  for $i = 1, ..., m$ **do**

    $c_{k+1}^i = q_{k+1}^i \cdot f_k$;

    $\partial_{\xi_k^i} f_k = D^2\varphi_k(v_k, q_k^i) + D\varphi_k\, w_k^i + D\partial_s\varphi_k\, q_k^i$;

  end

  $v_{k+1} = f_k - c_{k+1}^i\, q_{k+1}^i$;

  $\nabla_{\xi_{k+1}} f_k = \nabla_{\xi_k} f_k\, R_{k+1}^{-1}$;

  for $i, j = 1, ..., m$ **do**

    $p_{k+1}^{i,j} = a_{k+1}^{i,j} - q_{k+1}^l (\partial_{\xi_{k+1}^j} R_{k+1})^{(li)}$;

    $b_{k+1}^{i,j} = p_{k+1}^{i,j} \cdot f_k + q_{k+1}^i \cdot (\nabla_{\xi_{k+1}} f_k)^{:j}$;

  end

  for $i = 1, ..., m$ **do** $w_{k+1}^i = (\nabla_{\xi_{k+1}} f_k)^{:i} - b_{k+1}^{l,i}\, q_{k+1}^l - c_{k+1}^l\, p_{k+1}^{l,i}$;

  Save the scalar: $u_{k+1} = b_{k+1}^{i,i} + c_{k+1}^i\, g_{k+1}^i$;

  Advance the iteration: $x_{k+1} = \varphi(x_k)$;

  Evaluate: $D\varphi_{k+1}$, $D^2\varphi_{k+1}$, $D\partial_s\varphi_{k+1}$, χ_{k+2} , J_{k+1}, DJ_{k+1};

end

four nested for-loops, which implies that the brute-force vector-by-vector re-scaling would require $\mathcal{O}(n\,m^4)$ floating point operations. However, as pointed out in Section 5.3, this operation can also be completed in a component-by-component fashion. It means that one can alternatively compute n matrix products $(R_{k+1}^T)^{-1}\,\tilde{A}^i\,R_{k+1}^{-1}$, where $\tilde{A}^i$ denotes an $m \times m$ matrix containing i-th components of all m^2 instances of $\tilde{a}$. The double matrix-matrix product costs $\mathcal{O}(m^3)$ flops and thus the total cost of Line 11 is proportional to $\mathcal{O}(n\,m^3)$. The reader is referred to [162] for a more detailed analysis of the computational complexity of this part of the algorithm.

Lines 12-17 compute m upper-triangular derivatives of the R matrix. Each component requires evaluating one or two dot products, which implies that the cost of executing this chunk is $\mathcal{O}(n\,m^3)$. Here, we automatically obtain the SRB density gradient by evaluating the traces of all instances of $\partial_\xi R$. The simplified relation for g (Line 16) is a direct consequence of the measure conservation, which was obtained through the parametric differentiation of Eq. 5.17 using locally orthogonal coordinates (a complete derivation can be found in [160]). The leading order of the flop count of the code fragment involving Lines 18-23 is determined by Line 21. That line evaluates m parametric derivatives of f through the contraction of the Hessian and two other matrix-vector products. Per our discussion above, therefore, Line 21 requires $\mathcal{O}(n^3\,m)$ flops. The same estimate also applies to the algorithm part involving Lines 24-30. Here, the most expensive stage is Line 26, which evaluates m^2 instances of p. Note that for each instance of p, we compute a matrix-vector product, each requiring $\mathcal{O}(nm)$ algebraic operations. Note also the variable change (Line 24), computation of all scalars b (Line 27), and the update of w (Line 29) cost $\mathcal{O}(n\,m^2)$ each. The final chunk of this algorithm, Lines 31-32, evaluates a collection of nonlinear expressions and thus its complexity depends on the structure of the system.

Key point. Since $n \geq m$, the leading term of the total flop count of Algorithm 2 (excluding the nonlinear part) is proportional to $\mathcal{O}(n^3\,m^2)$. This estimate reflects the worst-case scenario for a general chaotic system. However, many real-world chaotic systems produced by spatial discretization of partial differential equations (PDEs) have a special structure. Popular discretization schemes, such as the finite element

method, generate systems with local dependencies. Each grid point is usually communicated only with the neighboring points. This implies that both the Jacobian and Hessian of φ are sparse arrays and have a banded structure. In such systems, therefore, the total flop count is linear with respect to the dimension of the system n. To conclude, the ultimate cost of approximating the sensitivity $d\langle J\rangle/ds$ using Algorithm 2 and data from N consecutive states along a typical trajectory is $\mathcal{O}(N\,n^3\,m^2)$. For physical systems, however, this estimate can be potentially reduced to $\mathcal{O}(N\,n\,m^3)$ in the presence of sparsity patterns arising due to a local discretization.

In terms of the storage, the largest arrays are the following: the Hessian $D^2\varphi$, m^2 vectors p, and $m^2/2$ vectors a. They collectively have n^3, $m^2\,n$ and $m^2\,n$ components, respectively. In the case of PDE-related systems with a sparse structure, the number of components to be stored is linear with respect to n. Moreover, in several physical dynamical systems, the dimension of the unstable manifold is significantly smaller than the system's dimension, i.e., $m/n \ll 1$ and $m^2 < n$ [18, 126, 21]. Thus, in the case of sparse physical systems, our algorithm requires storing two arrays no larger than n^2 entries and a few significantly smaller arrays. We encourage the reader to explore Supplementary Materials to see an example of efficient implementation of the Jacobian-vector and Hessian-vectors products for the n-dimensional solenoid map introduced in Section 5.5.

Note that in order to compute all required quantities at step $k+1$, we only need information from the previous one, i.e., k-th time step. No information from steps $k-2, k-3, ..., 0$ is required to advance the iteration from step k to $k+1$. This is a considerable advantage over other competitive methods such as the family of shadowing methods proposed in [18, 181, 129], which require storing large arrays from the entire history to assemble a huge global matrix. That matrix is a consequence of the minimization problem defined over the entire trajectory. Our method does not rely on global dependencies.

The desired sensitivity, $d\langle J\rangle/ds$, can be naïvely computed using the finite difference method (FD). However, the error associated with FD approximations combines both the statistical and Taylor series truncation errors. In practice, this means that

one might need even a few orders of magnitude more samples compared to Algorithm 2 to secure a similar error of the final output [158, 31]. We also highlight the fact that FD is not the best option for multi-query problems common in design optimization and uncertainty quantification. Problems of this type usually require estimating thousands of sensitivities of different observables with respect to large arrays of parameters. For FD, the total cost is obviously proportional to the number of parameters. This is not the case for the space-split algorithm, because its most expensive part, i.e., the part that is responsible for differentiating the SRB measure (Lines 7-17), does not depend on the parameter. Notice that g is computed only once per time step, regardless of the number of parameters.

5.5 Numerical results

The purpose of this section is to numerically analyze the convergence of Algorithm 2 and validate it using three different low-dimensional chaotic maps taken from the literature. In the first part, we shall consider the two-dimensional ($n = 2$) baker's map $\varphi : [0, 2\pi]^2 \to [0, 2\pi]^2$ [31],

$$\begin{aligned} x_{k+1}^{(1)} &= 2x_k^{(1)} + s^{(1)} \sin x_k^{(1)} + s^{(2)}/2 \sin x_k^{(1)} \sin 2x_k^{(2)} \bmod 2\pi \\ x_{k+1}^{(2)} &= x_k^{(2)}/2 + \pi\lfloor x_k/\pi \rfloor + s^{(3)}/2 \sin x_k^{(1)} \sin 2x_k^{(2)} + s^{(4)}/2 \sin 2x_k^{(2)} \bmod 2\pi, \end{aligned} \tag{5.38}$$

and the three-dimensional ($n = 3$) solenoid map [127, 186],

$$\begin{aligned} x_{k+1}^{(1)} &= 0.05x_k^{(1)} + 0.1\cos(8x_k^{(2)}) - 0.1\sin(5x_k^{(3)}) \\ x_{k+1}^{(2)} &= 2x_k^{(2)} + s\,(1 + x_k^{(1)})\,\sin(8x_k^{(2)}) \bmod 2\pi \\ x_{k+1}^{(3)} &= 3x_k^{(3)} + s\,(1 + x_k^{(1)})\,\cos(2x_k^{(3)}) \bmod 2\pi. \end{aligned} \tag{5.39}$$

The baker's map involves a set of four real-valued parameters, $\{s^{(1)}, s^{(2)}, s^{(3)}, s^{(4)}\}$. For moderately low parameter values, this map has one positive LE ($m = 1$) close to $\log 2$ and one negative LE. Eq. 5.38 is a mathematical representation of the kneading operation, in which a thin dough is stretched by the factor of 2 and then compressed

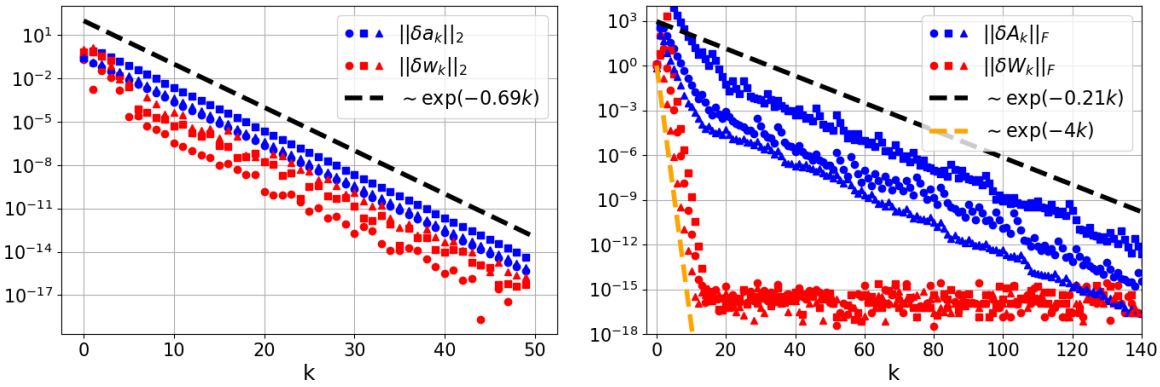

Figure 5-1: Convergence of the iterative formulas for a and w on the semi-logarithmic scale. We compute vector 2-norms and matrix Frobenius norms for the baker's map (left) and solenoid map (right), respectively. In both cases, the parameter values are randomly chosen from the uniform distribution $[0, 0.2]$. The dashed lines represent a reference exponential function. Each marker represents a distinct randomly chosen trajectory.

by the same factor. This stretching–compressing process is perturbed in two directions through the sine functions. The family of baker's maps serve as deterministic models of diffusion processes and are widely used in statistical mechanics [60]. The second map, in Eq. 5.39, is parameterized by a single real-valued parameter s. It was constructed in [127] by adding one additional expanding rotation and extra interaction term between contracting and expanding directions of the Smale-Williams map used in the modeling of oscillating circuits [93]. If s is moderately low, this map has two positive LEs ($m = 2$), with values close to $\log 2$ and $\log 3$, and a negative one. Therefore, unstable manifolds are geometrically represented by smooth curves immersed in $\mathbb{R}^2$ (baker's map) and surfaces immersed in $\mathbb{R}^3$ (solenoid map).

Figure 5.5 presents convergence plots of the recursive formulas for a and w, which are key ingredients of Algorithm 2. We perform this test by randomly choosing two different initial conditions w_0^i and $a_0^{i,j}$, $i, j = 1, ..., m$, running two independent sets of recursions, and computing the norm of the resulting difference vectors (as defined in Eq. 5.24 and Eq. 5.34) as a function of time k. We also randomly choose parameter values and repeat this test three times by following three randomly chosen trajectories.

We observe that both the acceleration vectors and parametric derivatives of the tangent solution obtained through the recursions derived in Section 5.3 converge ex-

ponentially fast, which confirms our analytical predictions. However, the iteration for the latter quantity tends to converge faster. For the 3D map, we observe a significantly larger convergence rate of the quantity represented by red markers. The corresponding errors reach the machine-precision values in only 10–12 iterations, after which the round-off error starts to dominate. This test also indicates that the rate of convergence may vary from system to system. We also notice a significant peak right after the beginning of the recursion in the case of the solenoid map. This is a consequence of the randomly chosen initial condition x_0 that is likely to be located beyond the attractor given its complex geometry [127].

We now turn our attention to a different "convergence" aspect of the S3 algorithm. Indeed, the overall accuracy of our method depends on the amount of data used in the ergodic-averaging process of the recursively generated time series. However, apart from the trajectory length N, we also have freedom in choosing the number of series terms K. To better understand the influence of that parameter on the error, we fix N and measure the accuracy of the output of Algorithm 2 by changing K. The error is defined as the ratio $|\tilde{x}_{S3} - \tilde{x}_{FD}|/|\tilde{x}_{FD}|$, where $\tilde{x}_{S3}$ and $\tilde{x}_{FD}$ denote the sensitivity approximations obtained using the S3 algorithm and finite differences, respectively. Figure 5-2 shows the relation between error and K for both maps. Each curve plotted in that figure was generated for a single trajectory, i.e, the result for K^* depends on the results for $K < K^*$. A few important conclusions follow from Figure 5-2. First, by fixing N, one can distinguish two dramatically different regimes of the error function. For low values of K, the error rapidly decays at the approximately exponential rate. As K gets larger, the S3 accuracy does not improve and the error might even increase. Second, by extending the trajectory, i.e., by increasing N, the qualitative behavior of the error function does not change, but the particular error values decrease. These two critical observations are consistent with the discussion associated with Eq. 5.37. Recall that the convergence of the S3 algorithm for $K \to \infty$ is guaranteed only if the particular ergodic averages (Lebesgue integrals) are exact. The exponential decay of correlations is manifested in our numerical experiment only to some extent, i.e., when K is small. Indeed, we observe that the exponentially decaying part of the

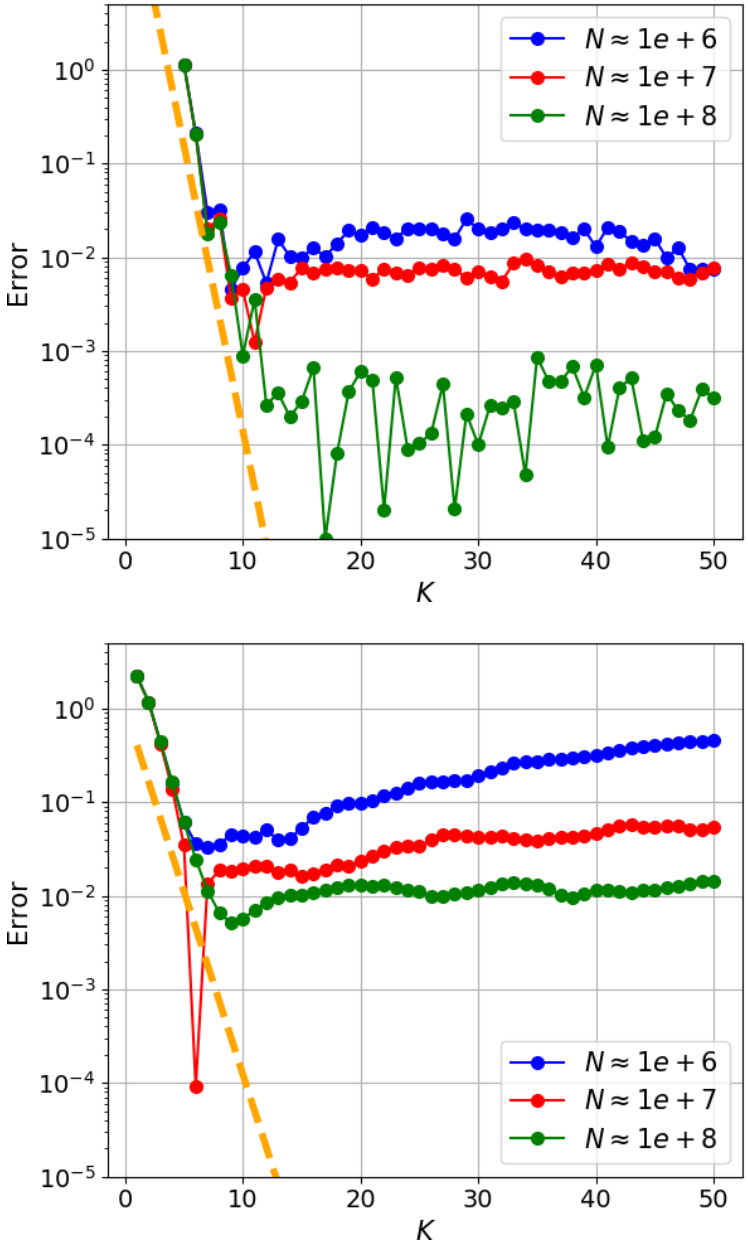

Figure 5-2: Relation between the error and K with a fixed trajectory length N. In this analysis, we tested Map 5.38 with $s = s^{(1)} = s^{(2)} = 0.6$, $s^{(3)} = s^{(4)} = 0$, $J = \cos(4\,x^{(1)})$ (left) and Map 5.39 with $s = -0.05$, $J = \sin(x^{(0)})\,\cos(4\,x^{(1)})\,x^{(2)}$. The error function represents a relative error computed with respect to the central finite difference approximation with grid size $\delta s = 0.01$. The approximations of long-time averages needed for finite differences were obtained using 10^{10} samples each. Dashed lines represent reference exponential functions.

error function extends as we increase N. A practical takeaway is the following. Let K^* be a critical value of K that indicates the transition between the two regimes of the error function. Notice that the value of K^* may vary between different systems. Since the cost of evaluating Line 5 of Algorithm 2 is negligible compared to the total cost, a good strategy is to compute and save several approximations of the unstable contribution along a single trajectory and choose the one that significantly breaks the converging trend.

We are also interested in the behavior of the error function with respect to the trajectory length N for a fixed K. Figure 5-3 illustrates that relation for both maps. All plots in Figure 5-3 clearly indicate that our algorithm behaves as a typical Monte Carlo procedure, i.e., the error approximately scales as $\mathcal{O}(1/\sqrt{N})$. This is consistent with the analytical predictions and our argument that the $\log\log N$ term has little impact on the convergence rate for large values of N. The relative error computed for the 3D map at $s = 0.05$ does not drop below the value 0.03 as we approach $N = 10^8$, which suggests that an individual calibration of K for that particular parameter value could be done. Notice also that we respectively need $N = \mathcal{O}(10^6)$ and $N = \mathcal{O}(10^8)$ samples to secure the relative error $\mathcal{O}(10^{-2})$, which means that the value of C from Ineq. 5.37 may significantly vary from system to system.

Given the above convergence test results, we shall now apply Algorithm 2 to Maps 5.38–5.39 for a wide range of parameters. In this experiment, we fix the cut-off threshold to $T = 150$, which will guarantee that the trajectory-following recursions provide us machine-precision quantities that will be used in ergodic-averaging. For both maps, we also fix $K = 11$, which is an average of the optimal values estimated in our convergence study. For every parameter value from the chosen interval, the S3 algorithm will run over $N = 10^6$ or $N = 10^7$ time steps. We validate all S3 outputs by comparing them against the central finite difference approximations (FD). The latter are obtained by computing long-time averages of a chosen objective function, fitting a polynomial, and applying the FD formula. Sensitivity analysis results for Map 5.38 and Map 5.39 are included in Figure 5-4 and Figure 5-5, respectively. Other technical details are included in the captions of these two figures.

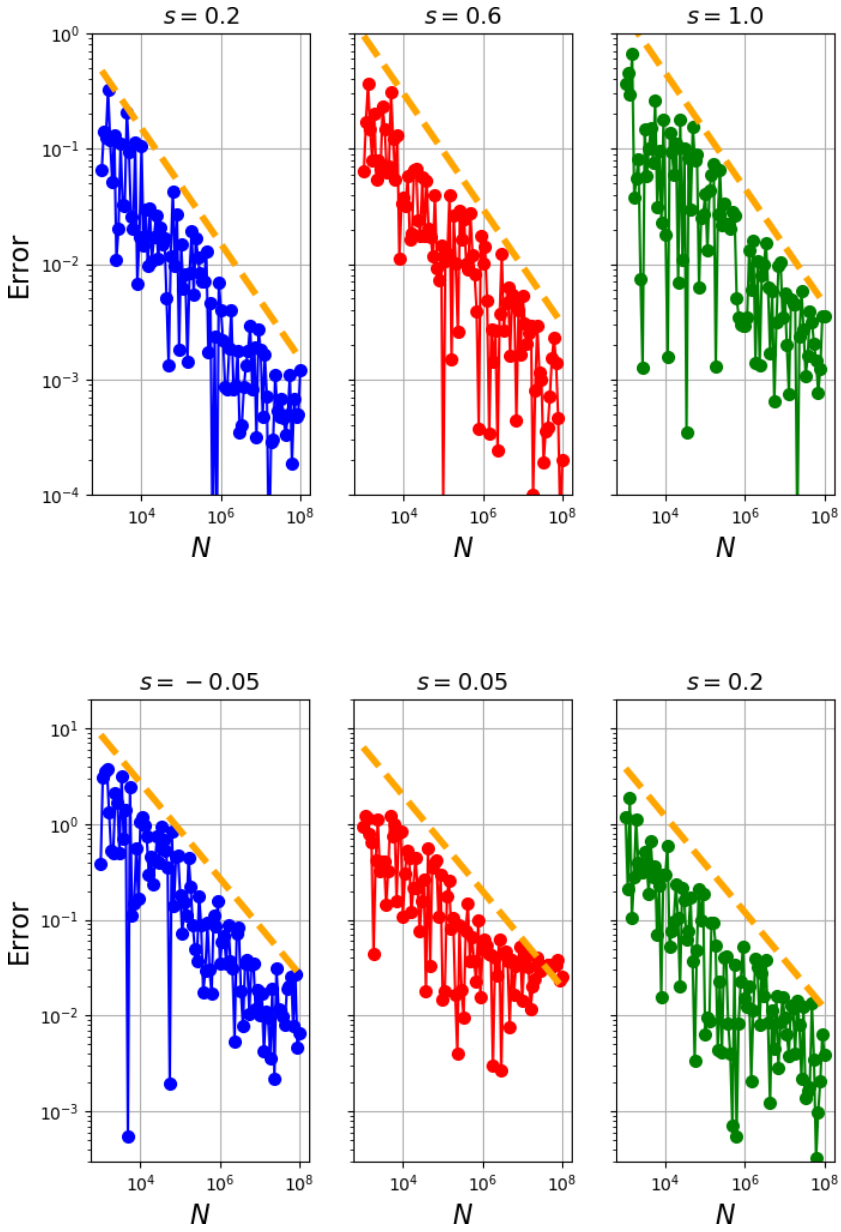

Figure 5-3: Relation between the error and trajectory length N for a fixed K. The top row represents Map 5.38 with $s = s^{(1)} = s^{(2)}$, $s^{(3)} = s^{(4)} = 0$ and fixed $K = 14$, while the bottom row represents Map 5.39 with $K = 8$. The objective functions and the way the error is computed are identical to those in Figure 5-2. The dashed lines are reference functions that are proportional to $1/\sqrt{N}$.

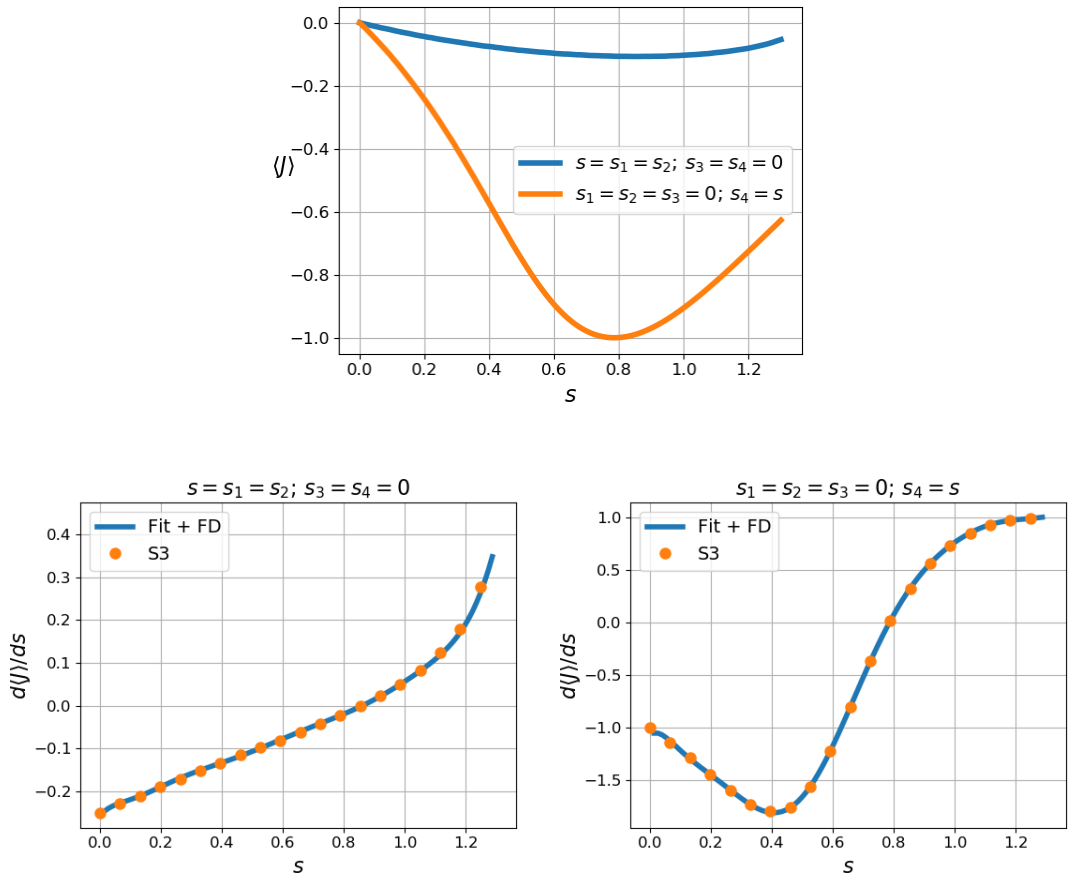

Figure 5-4: Sensitivity analysis of Map 5.38 for a wide range of parameters. Upper plot: long-time averages of $J = \cos(4\,x^{(1)})$ versus two sets of parameters. The two curves were created using 100 equally-spaced data points. For each grid point, a random trajectory of length $N = 10^8$ was computed. Lower plots: comparison of sensitivities computed in two different ways: through central finite differences (FD) and S3 algorithm using $N = 10^6$ samples. The FD scheme was applied to an 11th-order polynomial fit of the two curves in the top plot.

We observe that Algorithm 2 provides accurate approximations of the linear response for a wide range of parameters for both 2D and 3D maps. The curve representing FD approximations crosses all orange bullets despite the fact the S3 estimates were generated for a fixed K and N. Results of this experiment indicate that, in this particular case, there is no need to calibrate K and N for every instance of the input parameter s. If the parameter range is wider than the ones used in this study, a further calibration might be necessary to preserve the desired accuracy.

In the second part of this section, we apply Algorithm 2 to a problem that is more challenging than the two above. Thus, we shall consider a modified version of

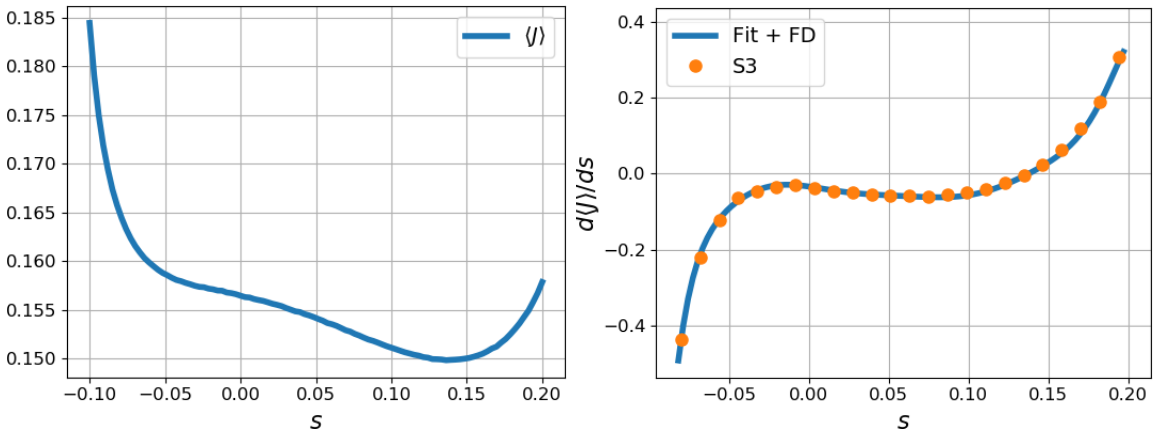

Figure 5-5: Sensitivity analysis of Map 5.39 for a wide range of parameters. Left plot: long-time averages of $J = \sin(x^{(0)})\cos(4\,x^{(1)})\,x^{(2)}$ versus s. The curve was created using 100 equally-spaced data points. For each grid point, a random trajectory of length $N = 10^8$ was computed. Right plot: comparison of sensitivities computed in two different ways: through central finite differences (FD) and S3 algorithm using $N = 10^7$ samples. The FD scheme was applied to a 11th-order polynomial fit of the curve in the left plot.

the n-dimensional solenoid map that was studied in [127]. We modify its original version by setting different expansion coefficients in phase space. The new map has the following form,

$$\begin{aligned} x_{k+1}^{(1)} &= 0.05\,x_k^{(1)} + s + 0.1\sum_{j=2}^{n}\cos(5\,x_k^{(j)}) \\ x_{k+1}^{(j)} &= \left(2+\frac{j-1}{n-1}\right)x_k^{(j)} + s\,(1+x_k^{(1)})\,\sin(2\,x_k^{(j)}) \bmod 2\pi, \quad j=2,...,n, \end{aligned} \tag{5.40}$$

where $s \in \mathbb{R}$, and $x_k \in \mathbb{R}\times[0,2\pi]^{n-1}$ for all $k \in \mathbb{Z}$. This is an n-dimensional extension of the classical solenoid map and was constructed by adding additional expansive directions. It can be shown that for a small s, Map 5.40 has one negative and $m = n-1$ positive Lyapunov exponents. In our experiments, we set $n = 12$, which means that the stable subspace consists of simple 1D curves immersed in $\mathbb{R}^{12}$. The sensitivities will be computed using the objective function from the study of the original map, i.e.,

$$J = (x^{(1)})^3 + 0.005\sum_{j=2}^{12}\sin(x^{(j)}-\pi).$$

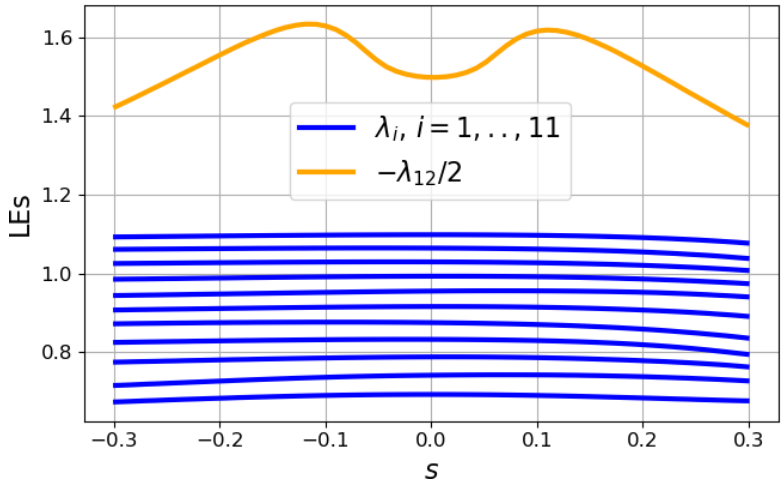

Figure 5-6: Lyapunov exponents of the 12-dimensional variant of Map 5.40. The spectrum was computed at 96 equally-spaced parameter values. For each s, a trajectory of length $N = 10^7$ was computed. The only negative LE, λ_{12}, was re-scaled by the factor $-1/2$ only for demonstration purposes.

Before we compute the linear response, we first investigate the Lyapunov spectrum of the modified map at $s \in [-0.3, 0.3]$. Recall that an integral component of the S3 algorithm is the recursive computation of orthogonal bases of unstable manifolds. Thus, by taking the ergodic-average of the logarithms of the diagonals of upper-triangular matrices R, we retrieve the positive part of the LE spectrum. We can modify Lines 7-9 of Algorithm 2 by adding extra $n-m$ tangent equations to retrieve the full Lyapunov spectrum. This method of approximating the spectrum was proposed in [15], which we use to plot all 12 exponents in Figure 5-6. We observe that the LE spectrum of Map 5.40 consists of 11 positive exponents scattered approximately between $\log 2$ and $\log 3$ and a negative exponent with the value oscillating around -3. The dimensions of the tangent subspaces are fixed over the entire parametric domain.

By setting $n = 12$, $m = 11$, and $N = 10^7$, $T = 150$, we run Algorithm 2 to differentiate the long-time average of J. Using the strategy described above, we store several approximations of the unstable contribution in our simulation, each corresponding to a different value of K. Figure 5-7 illustrates the statistics of J at $s \in [-0.3, 0.3]$ and corresponding sensitivities. The latter are approximated using central finite differences and S3 with $K = 2$, analogously to the previous numerical

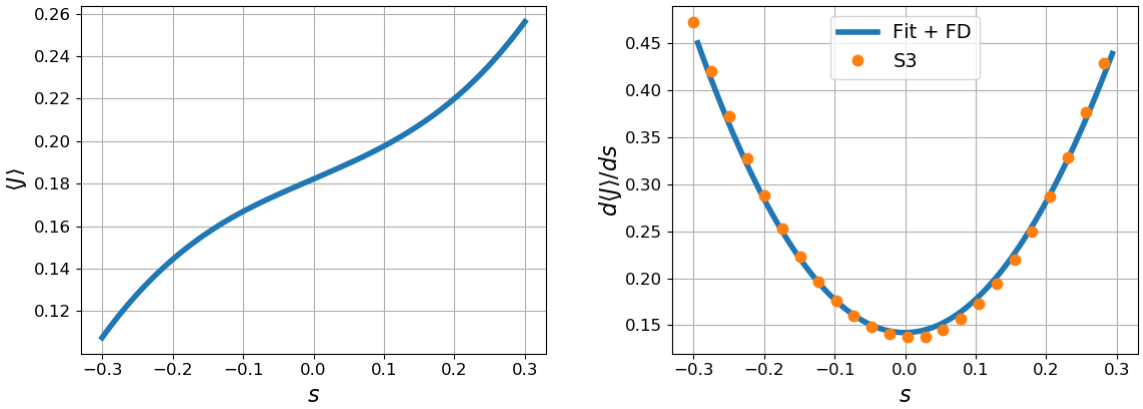

Figure 5-7: Sensitivity analysis of the 12-dimensional variant of Map 5.40. Left: long-time averages of J computed at 96 equally-spaced parameter values. For each s, a trajectory of length $N = 10^9$ was computed. Right: approximation of sensitivities using a 5th-order polynomial fit with central finite differences ($\delta s = 0.6/95$), and the S3 algorithm with $N = 10^7$, $T = 150$, $K = 2$.

experiments. Although K is fixed everywhere in the parametric space in Figure 5-7, the maximum relative error of S3 approximations is only about 5%. The largest error is located around $s \in [0, 0.1]$, which could be further reduced by adding more samples or tuning the value of K. This example proves the correctness of the derived linear response algorithm summarized in Algorithm 2, which opens the door for potential extensions to continuous-time chaos.

5.6 Summary

It is generally difficult to accurately estimate sensitivities of chaotic dynamical systems. Due to the butterfly effect, the direct simulation of state perturbations is impractical. Several numerical methods have been proposed to compute the sensitivity of chaos, but most of them suffer from at least one of the following common problems: exploding tangent solutions, unphysicality of shadowing trajectories, huge computational cost and storage requirements, complicated generalization.

Our new method for sensitivity analysis derives from Ruelle's rigorous linear response theory, which is regularized based the concept of perturbation space-splitting and partial integration along unstable manifolds [31]. Through the intuitive measure-

based parameterization of the unstable subspace [160, 162] and chain rule on smooth manifolds, we systematically derive a set of trajectory-following recursions for a collection of by-products arising due to the regularization of Ruelle's formula, and analytically show their exponential convergence. Similarly to the majority of methods that stem from the linear response theory, our method is formulated as a Monte Carlo procedure, which provably converges to the true solution approximately as $\mathcal{O}(1/\sqrt{N})$, where N is the trajectory length. The following list summarizes the main advantages of the space-split approach:

- Immunity to the ergodicity-breaking and unphysicality errors, which are common in shadowing methods,
- Immunity to the exponential growth of norms of tangent solutions, a.k.a. the butterfly effect phenomenon,
- Generalizability to n-dimensional systems, $n \in \mathbb{Z}^+$; the algorithm we propose is ready-to-use for discrete systems with an arbitrary number of positive Lyapuonv exponents m,
- Provable convergence for uniformly hyperbolic systems,
- Translatable to memory-efficient as-we-go Monte Carlo algorithms.

The major consequence of partial Lebesgue integration is the computation of directional derivatives of the ergodic measure that describes the statistical behavior in phase space. Recursive computation of this quantity, known as the SRB density gradient, requires solving a collection first- and second-order tangent equations. Indeed, this is the actual price for the regularized version of Ruelle's formula. Therefore, from the algorithmic perspective, we must perform a series of algebraic operations involving third-order tensors. We estimate that the total flop count is $\mathcal{O}(n^3\, m^2)$, which can be reduced to $\mathcal{O}(n\, m^3)$ for PDE-related systems with local dependencies and sparse structures.

While this chapter solely focuses on discrete systems, the proposed algorithm can be naturally extended to continuous-time (ODE) systems. The perturbation vector

splitting would need to incorporate the one-dimensional neutral subspace that is aligned with the flow. This will require some modifications of the proposed algorithm, as the time evolution vector will need to be differentiated along unstable manifolds as well. However, after that minor modification, the leading term of the algorithm's total flop count is expected to remain the same.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 6

Approximating the linear response of physical chaos

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 chapter 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.

The content of this chapter was published in *Nonlinear Dynamics* by *Springer Nature* [163]. It is reproduced with permission from *Springer Nature*.

6.1 Preliminaries

In this chapter, we investigate if 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 literate 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 [18, 126]. Moreover, a recent theoretical study in [128] 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 [85], turbulence [97], population dynamics [177], 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 perturba-

tion splitting proposed and numerically tested in [161]. 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 [15]. This cascade property was used in [161] 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.

6.2 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 chapter, we consider a parameterized n-dimensional *ergodic* flow,

$$\frac{dx}{dt} = f(x; s), \quad x(0) = x_0, \tag{6.1}$$

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 [7]. 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 $J : M \to \mathbb{R}$, our ultimate goal is to approximate the parametric derivative of the long-time average of J, defined as

$$\frac{d\langle J \rangle}{ds} := \frac{d}{ds} \lim_{T \to \infty} \frac{1}{T} \int_0^T J(x(t;s))\, dt, \tag{6.2}$$

where M denotes the n-dimensional manifold defined by Eq. 6.1. We assume J does not depend on s.

Motivated by the work of Ruelle [148, 150], the authors of [31, 28] proposed a new method, called the space-split sensitivity (S3), which regularizes Ruelle's series (see Eq. 1.5) for systems with one-dimensional unstable subspaces ($m = 1$). Based on its extension to general hyperbolic maps in proposed in Chapter 5, 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, thoroughly described in the previous chapters of this thesis, is to decompose the perturbation vector χ into three terms,

$$\chi = \chi^u + \chi^c + \chi^s = \left(\sum_{i=0}^{m} c^i\, q^i \right) + \left(c^0\, f \right) + \left(\chi - \sum_{i=0}^{m} c^i\, q^i - c^0\, f \right), \tag{6.3}$$

such that χ^u and χ^c strictly belong to the unstable and neutral/center subspaces, respectively. In this splitting, c^i, $i = 0, ..., m$ are some scalars that are differentiable on the unstable subspace defined by a local orthonormal basis q^i, $i = 1, ..., m$. 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. 6.3:

- 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, $i = 0, ..., m$ 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. 1.5, we conclude that

$$\sum_{t=0}^{\infty} \int_M D(J \circ \varphi^t) \cdot \chi^s \, d\mu = \int_M DJ \cdot v \, d\mu, \tag{6.4}$$

where

$$v_{k+1} = D\varphi_k \, v_k + \left(\chi_{k+1} - \sum_{i=0}^{m} c^i_{k+1} \, q^i_{k+1} - c^0_{k+1} \, f_{k+1} \right). \tag{6.5}$$

The subscript notation indicates the time step, i.e.,

$$f(x(k\Delta t)) := f_k,$$

assuming uniform time discretization. To solve Eq. 6.5, 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 Section 6.2.2. 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,

$$(f_{k+1} \cdot f_{k+1}) \, c^0_{k+1} = f_{k+1} \cdot \left(r_{k+1} - \sum_{i=1}^{m} c^i_{k+1} \, q^i_{k+1} \right), \tag{6.6}$$

$$c^i_{k+1} = q^i_{k+1} \cdot \left(r_{k+1} - c^0_{k+1}\, f_{k+1}\right), \;\; i = 1, ..., m. \tag{6.7}$$

Eq. 6.6–6.7 define a linear system with $m+1$ equations and $m+1$ unknowns (c^i, $i = 0, 1, ..., m$). The system's matrix involves an $m \times m$ identity block I, while its Schur complement can be expressed as follows,

$$S_{k+1} = I - \frac{Q^T_{k+1} f_{k+1} (Q^T_{k+1} f_{k+1})^T}{f_{k+1} \cdot f_{k+1}}, \tag{6.8}$$

where Q is a an $n \times m$ matrix containing an orthonormal basis of the unstable manifold, q^i, $i = 1, ..., m$. Thus, the coefficients c^i, $i = 1, ..., m$, stored in the array c are obtained by solving the following reduced system,

$$S_{k+1}\, c_{k+1} = Q^T_{k+1} \left(r_{k+1} - \frac{f_{k+1} \cdot r_{k+1}}{f_{k+1} \cdot f_{k+1}} f_{k+1}\right), \tag{6.9}$$

while c^0 is computed directly from Eq. 6.6. We conclude that the stable part of the linear response can be evaluated through the ergodic average of $DJ \cdot v$ (see Eq. 1.6), where v satisfies Eq. 6.5–6.7.

The next step is the neutral contribution, which involves the perturbation component that is parallel to f. Analogously to Eq. 1.7, we can expand

$$D(J \circ \varphi^t) \cdot \chi^c = D(J \circ \varphi^t) \cdot (c^0\, f) = c^0\, DJ_t \cdot (D\varphi_{t-1} ... D\varphi\, f)\,. \tag{6.10}$$

Applying the Taylor series expansion, we note that

$$f(\varphi(x)) = f(x) + Df(x)\,(\varphi(x) - x) + \mathcal{O}((\varphi(x) - x)^2), \tag{6.11}$$

and, analogously,

$$\varphi(x) = x + \Delta t\, Df(x) + \mathcal{O}(\Delta t^2). \tag{6.12}$$

By differentiating Eq. 6.12 and plugging it to Eq. 6.11, we notice that in the limit $\Delta t \to 0$ we retrieve the covariance property, which reads

$$f(\varphi(x)) = D\varphi(x)\, f(x). \tag{6.13}$$

This implies that the neutral part can be simplified to

$$\sum_{t=0}^{\infty} \int_M D(J \circ \varphi^t) \cdot \chi^c \, d\mu = \sum_{t=0}^{\infty} \int_M c^0 DJ_t \cdot f_t \, d\mu = \sum_{t=0}^{\infty} \int_M c^0_{-t} DJ \cdot f \, d\mu. \tag{6.14}$$

Eq. 6.14 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 $t \to \infty$ [36, 195], i.e.,

$$\left| \int_M (J \circ \varphi^t) \, h \, d\mu - \int_M J \, d\mu \int_M h \, d\mu \right| \leq C\delta^t \tag{6.15}$$

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. 1.6.

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 [161]

$$\begin{aligned} &\sum_{t=0}^{\infty} \int_M D(J \circ \varphi^t) \cdot \chi^u \, d\mu = \sum_{t=0}^{\infty} \sum_{i=0}^{m} \int_M c^i \partial_{q^i} (J \circ \varphi^t) \, d\mu \\ &= -\sum_{t=0}^{\infty} \sum_{i=1}^{m} \int_M (J \circ \varphi^t) \left(c^i \, g^i + b^{i,i} \right) \, d\mu, \end{aligned} \tag{6.16}$$

where

$$b^{i,j} := \partial_{q^j} c^i, \quad g^i := \frac{\partial_{q^i} \rho}{\rho}, \tag{6.17}$$

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. 6.16. The reader is referred to Chapter 3 and Chapter 5 for a detailed description of every step of this process and relevant numerical examples. The major implication of Eq. 6.16

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 been proposed in the aforementioned two chapters. That algorithm requires solving a collection of first- and second-order tangent equations, and was developed for discrete chaotic systems. In Appendix 6.6.1, 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 (Eq. 1.5) 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 [31] and computational (see Chapter 5) 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.

6.2.1 Numerical example: Lorenz 63

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

$$\frac{dx}{dt} = \sigma(y - x), \quad \frac{dy}{dt} = x(\rho - z) - y, \quad \frac{dz}{dt} = xy - \beta z, \tag{6.18}$$

which is one of the simplest chaotic flows. The original study of this model [110] demonstrated chaotic behavior at $\sigma = 10$, $\beta = 8/3$, $\rho \gtrapprox 24$. For this choice of parameters, the strange attractor has a characteristic butterfly-shaped structure as illustrated in Section 4.5. 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

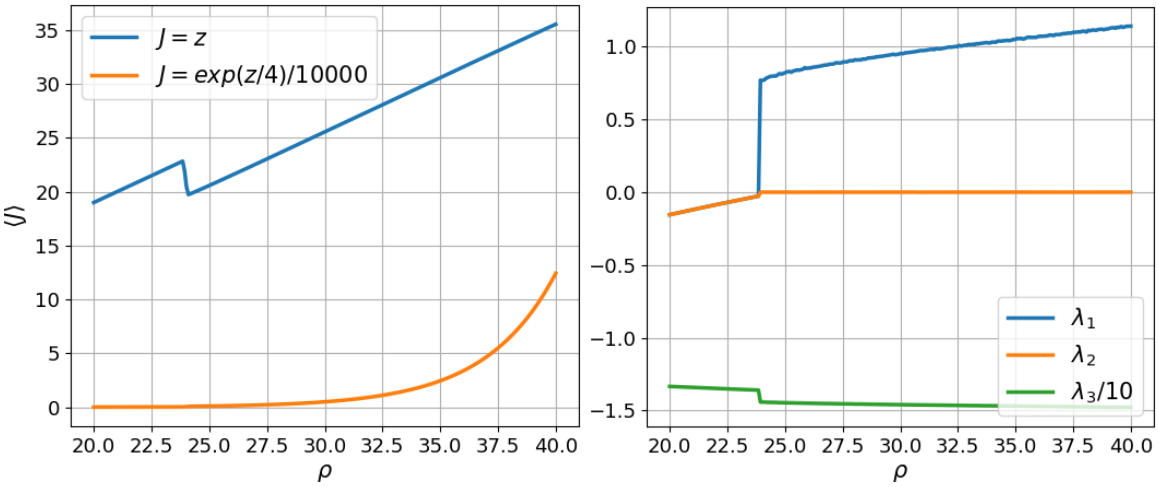

Figure 6-1: Long-time averages of two different objective functions (left) and Lyapunov exponents (right) versus the Rayleigh parameter ρ. Ergodic averages have been taken over $N\Delta t = 50,000,000$ and $N\Delta t = 5,000$ time units, respectively.

ρ using S3. In this section, ρ should not be confused with the SRB measure density. Figure 6-1 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 $\rho \gtrapprox 24$, 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. 6.18 in time, we used the second-order explicit Runge-Kutta with step size $\Delta t = 0.005$. As described in Appendix 6.6.1, 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 6.6.2. Fortunately, in the case of the Lorenz 63 system, $D^2 f(\cdot,\cdot)$ is constant, which significantly reduces the cost.

The S3 algorithm relies on several recursive formulas in the form of tangent equations. Earlier studies [31, 161] proved both analytically and numerically that these recursions converge exponentially fast in discrete hyperbolic systems. We numerically investigate if these results still apply to the Lorenz 63 flow. The upper plot of Figure

6-2 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 these quantities, we need only 50 time units. A similar behavior has been observed in the case of discrete systems [161]. 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 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 Figure 6-2 shows the decay of the relative error of the linear response approximation, which is computed with respect to the finite difference approximation of the slope of statistics generated in Figure 6-1. 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 [161], 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 Figure 6-1 at $\rho \in [25, 40]$. Figure 6-3 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

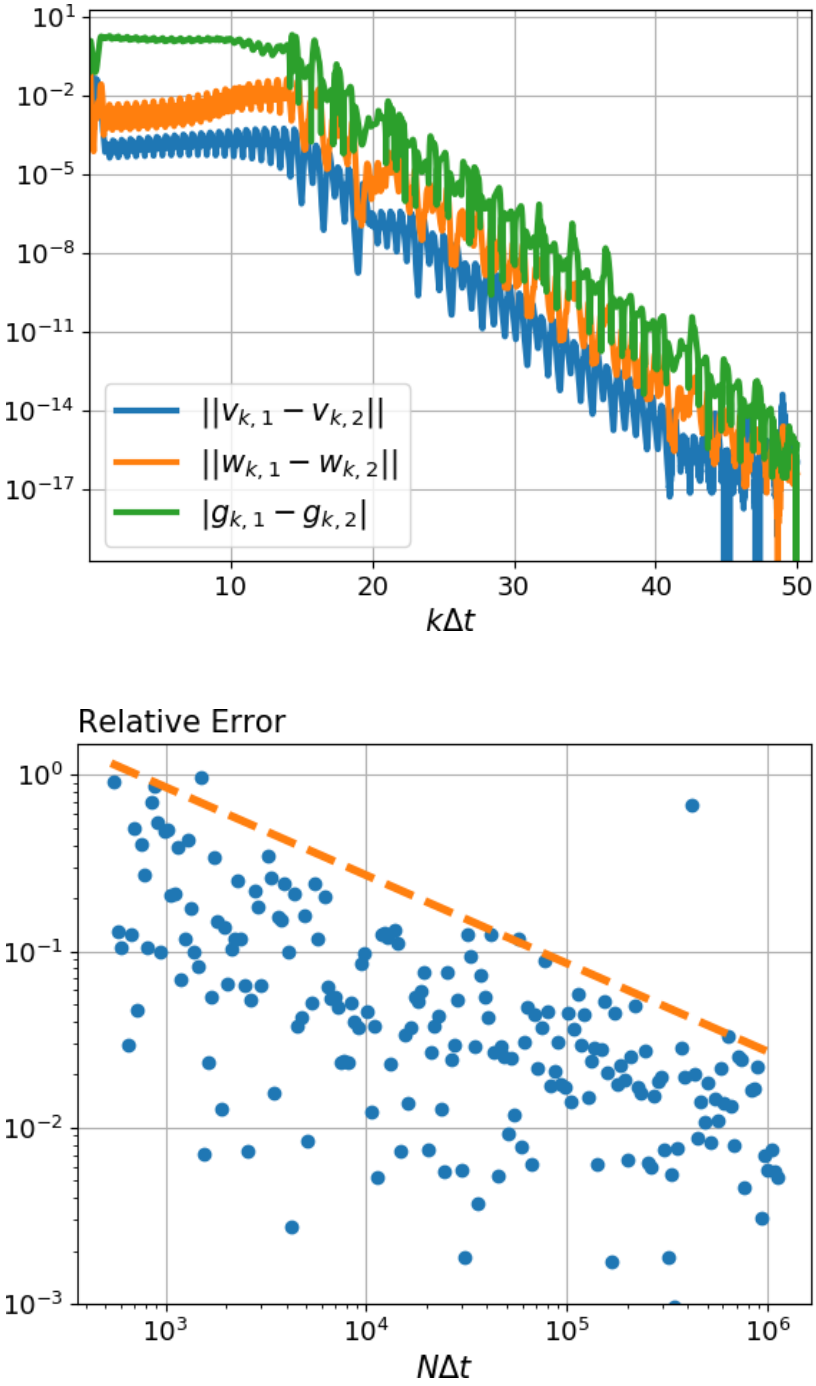

Figure 6-2: Upper: Relation of the norm/absolute value of the difference between quantities obtained along two different random orbits, labelled as 1 and 2, and time-averaging window $k\Delta t$. Lower: Relative error of the linear response approximation versus time-averaging window, computed for $J = z$ at $\rho = 28$. 200 independent simulations were run at a logarithmically uniform grid of $N\Delta t$. The dashed line represents a function $C/\sqrt{N\Delta t}$, $C > 0$.

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.

6.2.2 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.

Special treatment of the neutral component

We derived a numerical scheme based on the three-term linear splitting in Eq. 6.3. Notice that 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 Figure 6-4, 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 6.6.1, that minor conceptual adjustment requires major modifications of the “discrete" version of S3.

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 chapter

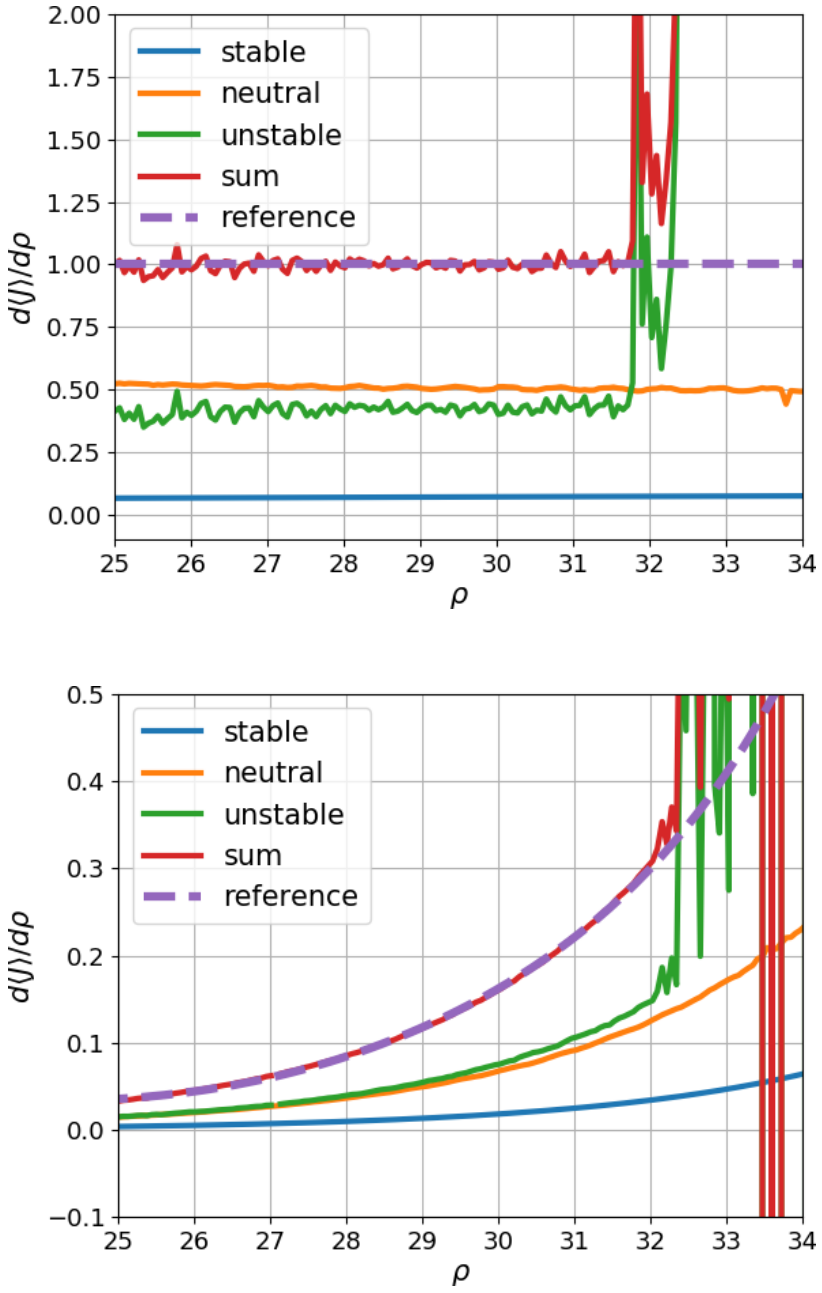

Figure 6-3: Output of Algorithm 2 generated for $J = z$ (upper) and $J = \exp(x/4)/10000$ (lower) at 144 values of ρ distributed uniformly. Each simulation was run for $N\Delta t = 1,000,000$ time units. The reference solution (dashed curve) was obtained using central finite differences and data shown in Figure 6-1. Before differentiation, we interpolated the data using first- and sixth-order polynomial fits, respectively.

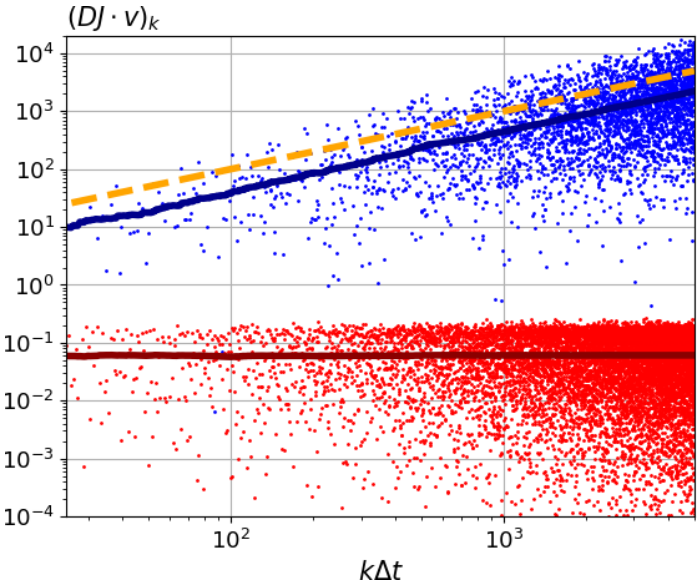

Figure 6-4: Discrete values of the stable integrand $DJ \cdot v$ computed using the S3 version described in Section 6.2 (red) and its "discrete" counterpart from [161] (blue). This simulation was performed for Lorenz 63 at $\rho = 28$. The solid lines represent the standard deviations of $DJ \cdot v$ collected from the beginning of the simulation until kth step. The dashed line represents a linear function.

or the shadowing-based variant proposed in [127], 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 Figure 6-5, 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 hyperbolicty verification proposed by Kuptsov in [89]. 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 [169]. If the statistical distribution of α is not strictly separated from the origin, i.e., the corresponding PDF has non-zero 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 center subspaces. Several numerical examples presented in [89] 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 [122] 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 [159], 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. 1.5 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 [89, 169] 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 [126]. 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.

Implementation and cost

We shall now comment on practical aspects of the full linear response algorithm, which is described in Appendix 6.6.1. 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

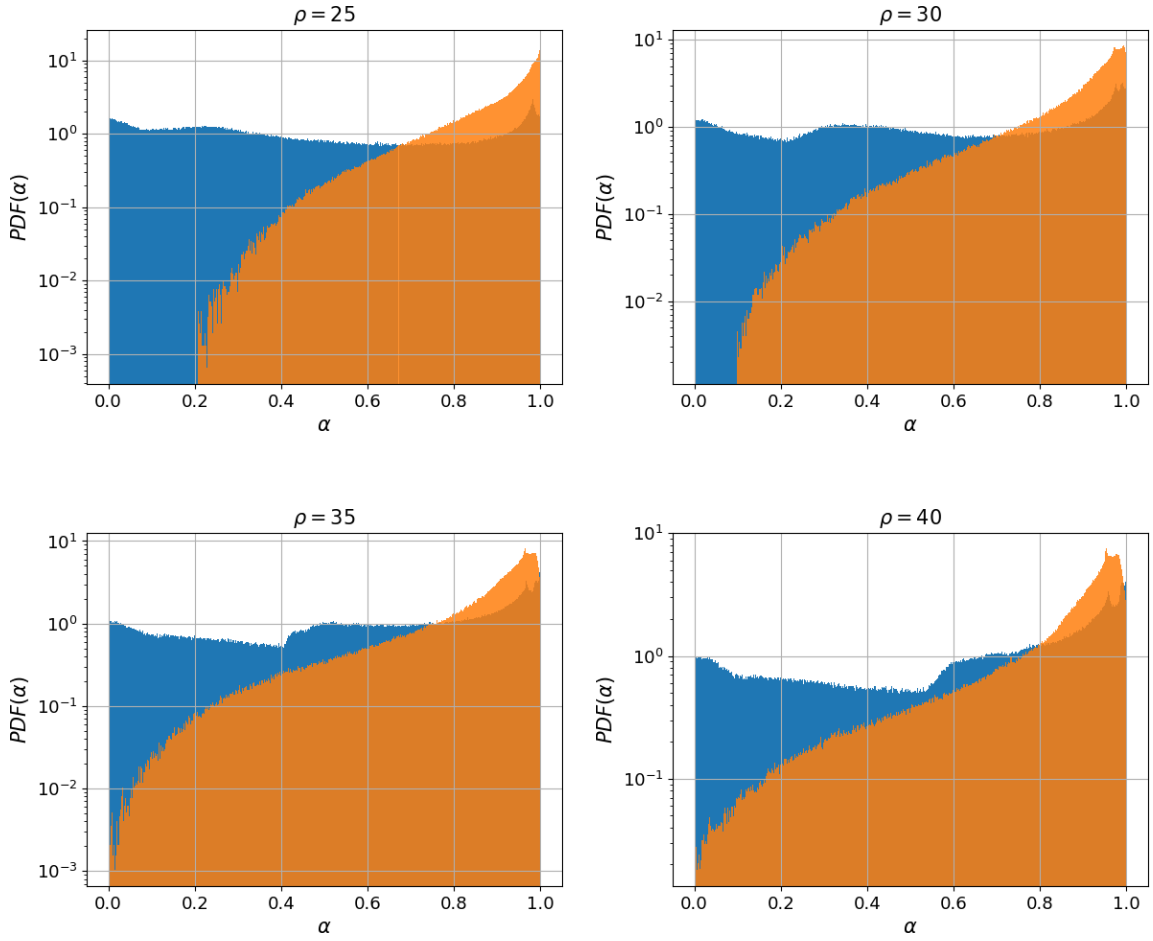

Figure 6-5: Distribution of the normalized measure α between unstable/center subspaces (blue PDF) and unstable-center/stable subspaces (orange PDF). They have been computed along a random trajectory of the Lorenz 63 system for 5000 time units. To increase the accuracy of PDFs, we used the fourth-order Runge-Kutta time integrator with $\Delta t = 0.005$.

through QR factorization that is needed to find a new basis of the subspace (matrix Q) and the Jacobian 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 [15]. 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 is 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 6.6.1, 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 $\mathcal{O}(n\,m^3)$ flops. This might pose a serious challenge for systems with hundreds of unstable modes, such as 3D turbulence models.

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 [92], 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 [159], 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 6-3 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 chapter.

6.3 Unstable contribution: can we neglect that term?

As we pointed out in Section 6.2.2, 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.

6.3.1 Empirical evidence of decaying components of the measure gradient

Recall that the unstable contribution depends on the objective function, splitting coefficients and their derivatives, and the SRB measure gradient. While the splitting coefficients directly follow from the linear splitting of Ruelle's formula and their values are chosen such that the stable tangent is bounded, the measure gradient is an inherent property of the system and appears in the unstable contribution only. The g vector represents the measure change in m orthogonal directions of the local unstable subspace. These directions, stored in the Q matrix, indicate how a trajectory of the unperturbed system deforms in time. We also acknowledge that the computation of Q is an integral part of the S3 procedure (see Appendix 6.6.1). In that algorithm, the columns of Q are sorted from the most expansive ($i = 1$) to the least expansive

$(i = m)$ direction.

In this analysis, we leverage our *observation* on monotonically increasing norms of consecutive components of the SRB measure gradients. In particular, we observe that in several systems with multidimensional expansive tangent subspaces

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

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

$$\|h\|_2 := \sqrt{\int_M h^2 \, d\mu} \tag{6.19}$$

for any scalar function $h \in L^2(\mu)$. Differences in norm between different components of g might be even orders of magnitude large if the positive Lyapunov exponents are not close to each other. The major conclusion that follows from our observation is that the slope of SRB measure computed in the most expansive direction is usually small compared to the remaining directions of the unstable manifolds everywhere on the attractor. In other words, larger expansion rates lead to the dilution of measure, which consequently decreases the corresponding measure slopes.

We also emphasize the fact that specific directions of differentiation, indicated by q^i, $i = 2, ..., m$, depend on the type of manifold parameterization. The only exception is q^1, which coincides with the leading covariant direction of the tangent manifold. Note also that the span of the subset of k leading Lyapunov vectors, $\text{span}\{q^1, q^2, ..., q^k\}$, $k \leq m$, does not depend on the parameterization either. By definition, first k components of g represent rates of measure change within the most expansive subset. Thus, the *cumulative* measure change within that subset, constituted by $g^1, g^2, ..., g^k$, $k \leq m$, neither depends on the parameterization. In other words, the unstable divergence of measure (or any other smooth observable) is invariant to the coordinate transformation, although some ingredients might change.

To illustrate our critical observations, we will focus on the following n-dimensional

chaotic map $\varphi : [0, 2\pi]^n \to [0, 2\pi]^n$ defined as

$$x^i_{k+1} = 2\, x^i_k + s\, \sin(x^{i+1}_k - x^i_k) + t\, \sin(x^i_k) \bmod 2\pi, \;\; i = 1, .., n, \tag{6.20}$$

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* [158], and therefore we shall refer to φ defined by Eq. 6.20 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 6.6.1) and plot both $|g^1|$ and $|g^2|$ on $[0, 2\pi]^2$. These results are illustrated in Figure 6-6. 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 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 $|s| \to 0$.

In Figure 6-7, 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 6-7 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 [159]. 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 Section 6.2.2.

Another example represents higher-dimensional continuous-time chaos. We will consider 40-dimensional Lorenz 96 oscillator with constant forcing of value 8. The reader is referred to Section 6.4.2 for a detailed description of the underlying ODE system and its physical interpretation. The system we consider here has 13 positive LEs that are approximately uniformly scattered between 0 and 1.8 (see Figure 6-11). Figure 6-8 illustrates PDFs of individual components of the SRB measure gradient approximated using data collected along a million time steps.

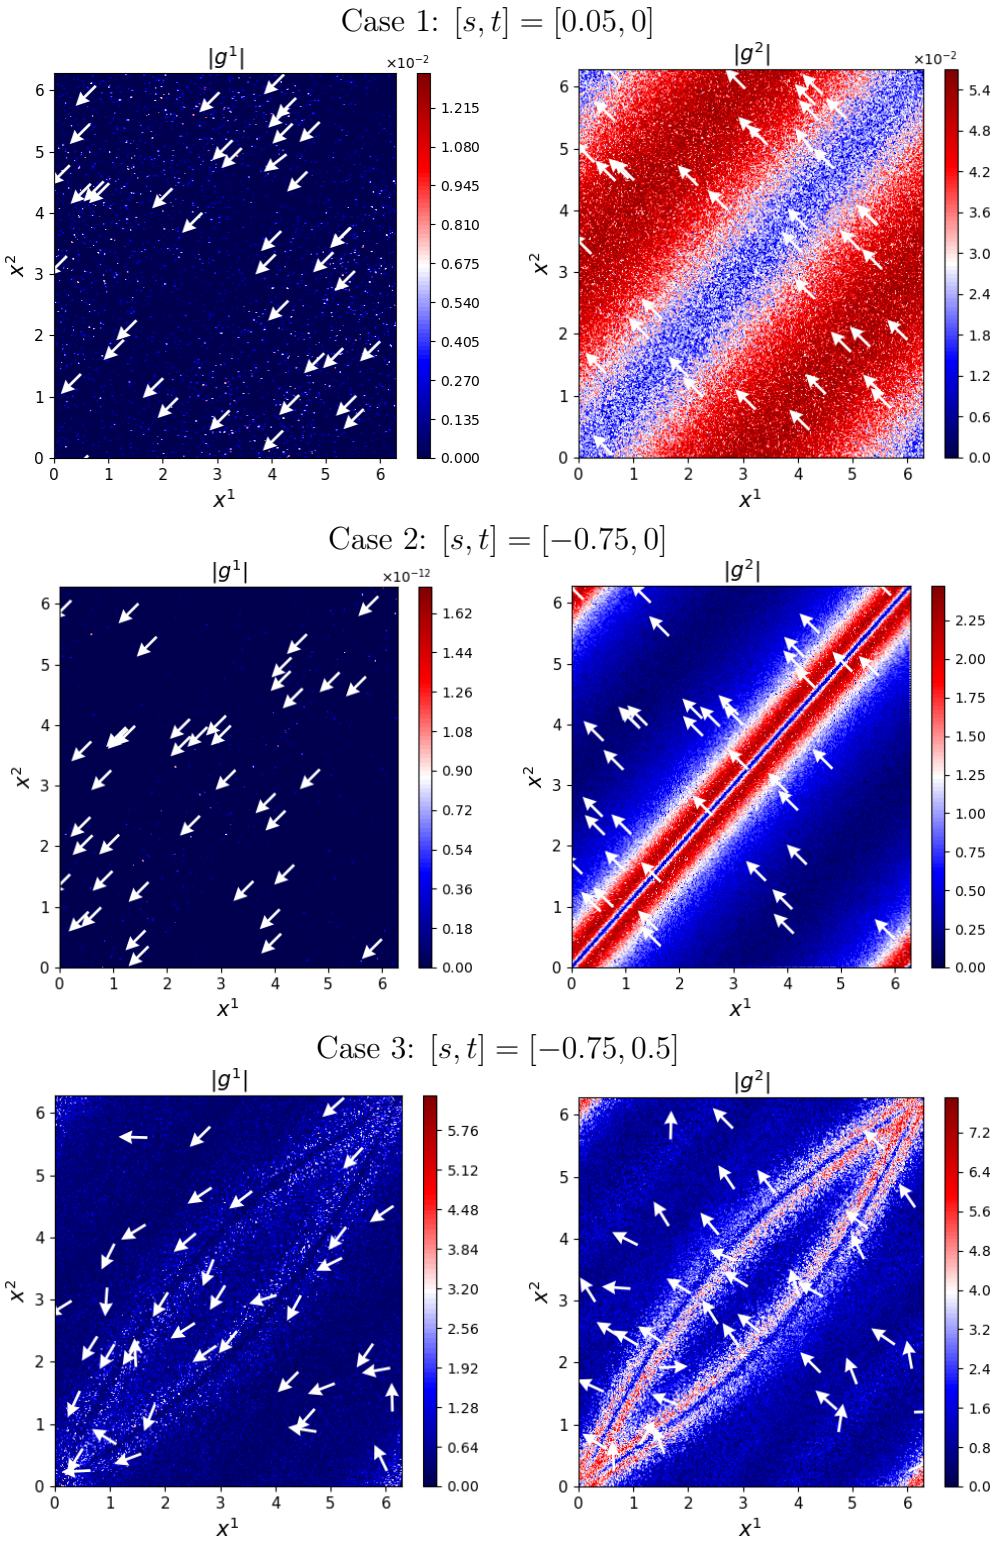

Figure 6-6: Magnitude of both components of the SRB density gradient g of the two-dimensional coupled sawtooth map with two positive LEs. White arrows respectively represent q^1 and q^2, which indicate local directions of differentiation. They are plotted every 5000 time steps. For each case, a trajectory of length $N = 3 \cdot 10^5$ was generated.

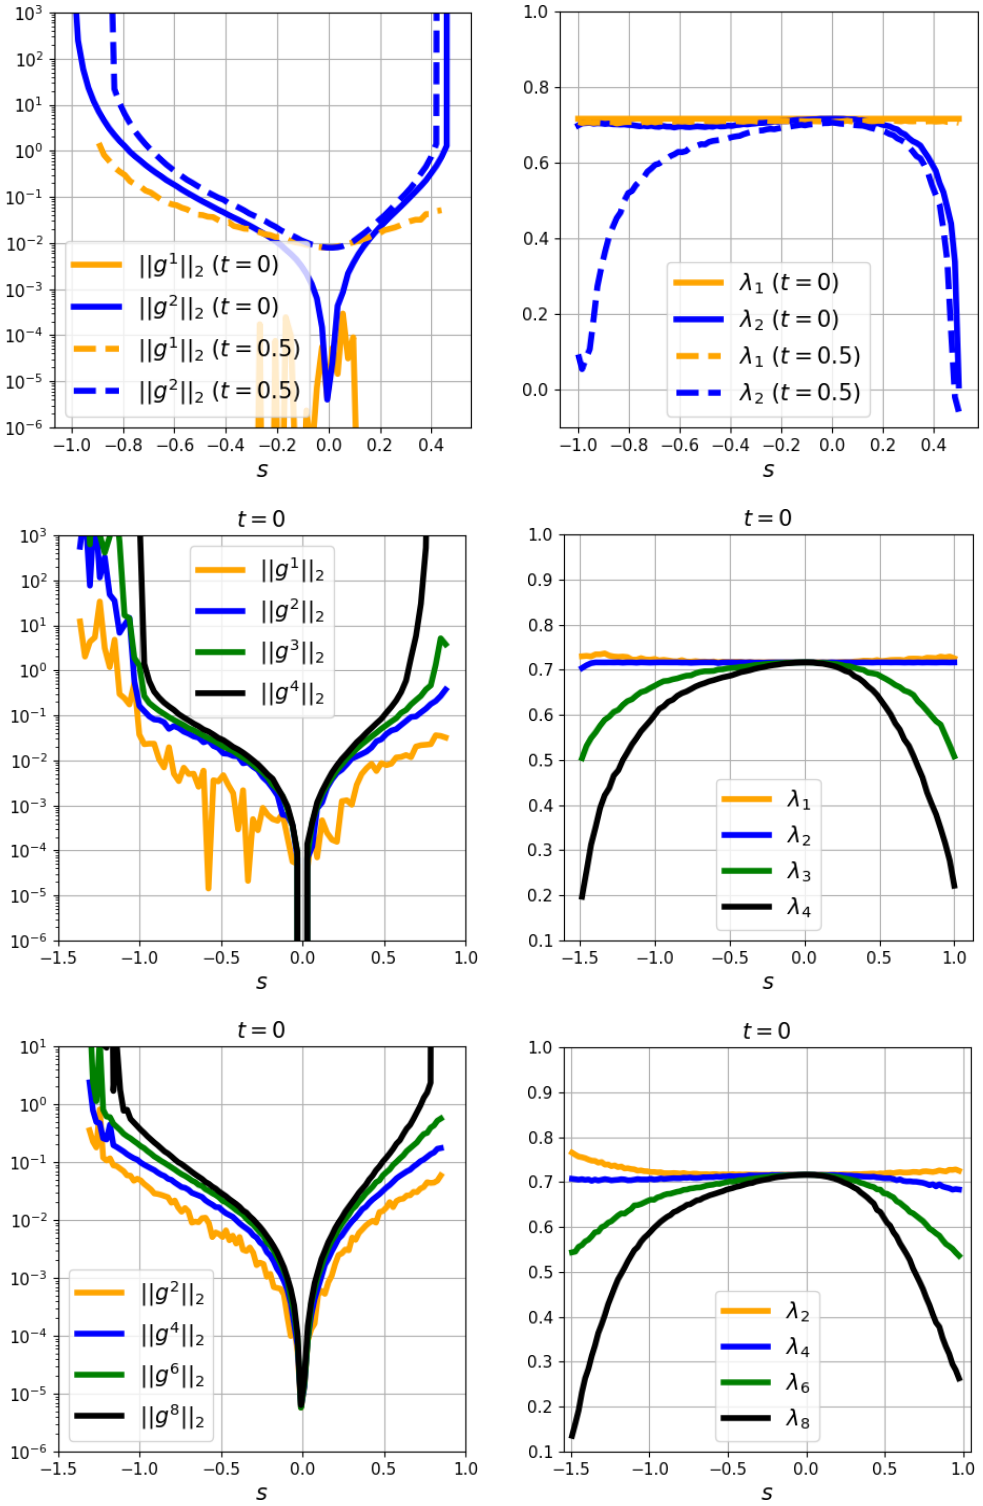

Figure 6-7: L^2 norm of the SRB density gradient and Lyapunov exponents of the 2D ($n = 2$; top row), 4D ($n = 4$; middle row), and 8D ($n = 8$; bottom row) variant of the coupled sawtooth map. All quantities were computed on a uniform grid of 100 values of the coupling parameter s. For each parameter, a trajectory of length $N = 3 \cdot 10^4$ was generated.

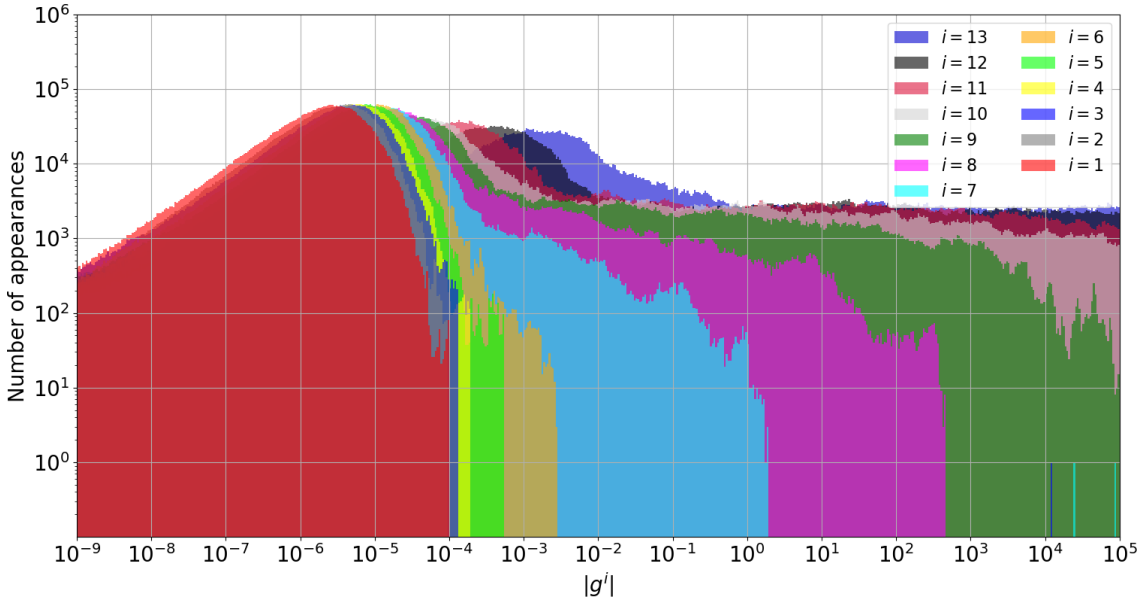

Figure 6-8: Distribution of individual components of the SRB measure gradient of Lorenz 96. SRB measure was differentiated using Algorithm 2 for $N = 10^6$ time steps. ODEs were discretized in time using the second-order Runge-Kutta with time step size 0.001.

We observe that the distribution tail monotonically gets longer as the component index increases. This behavior is analogous to the monotonically increasing L_2 norms of the coupled sawtooth map. Here, only the first three leading components are likely to have finite L_2 norms. If $i > 6$, the measure gradient is not even Lebesgue-integrable in the corresponding directions. We refer the reader to Chapter 4 for further details on assessing the differentiability of observables. In the next few paragraphs, we will leverage the fact the leading (or a subset of leading) components of the measure gradient might be significantly smaller compared to the remaining ones. In particular, we will show that this feature could dramatically reduce the magnitude of the unstable contribution.

6.3.2 Impact of decaying components of the measure gradient

Let us now revisit the analytical form of the unstable contribution,

$$U := \sum_{k=0}^{\infty} \int_M J \circ \varphi^k \, (c \cdot g + b) \, d\mu, \tag{6.21}$$

where c and g are m-dimensional vectors, representing the splitting coefficients and measure gradient, respectively. The scalar b represents the sum of derivatives of c. Thus, U is in fact an infinite series of time correlations between J and a linear combination of all components of g, i.e., $d := c \cdot g + b$. Let us now consider a well-behaved objective function $J : M \to \mathbb{R}$, 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.,

$$J = \nabla_\xi \cdot Z. \tag{6.22}$$

The nabla operator ∇_ξ contains parametric differentiation operators in the orthonormalized measure-based coordinate system as introduced in Chapter 3. Recall that the column space of the chart gradient is orthonormal. Note that the above representation of J is not unique. If M is one-dimensional, Z can be any function from the family of antiderivatives of J that differ by a constant. However, if the dimension of M is at least two, one can find Z with various types of behavior along different phase space directions.

By combining Eq. 6.21 and 6.22, using the measure preservation property and integrating by parts, we conclude that

$$\begin{aligned} U =& \sum_{k=0}^{\infty} \int_M J \circ \varphi^k \, d \, d\mu, \\ =& \sum_{k=0}^{\infty} \int_M J \, d \circ \varphi^{-k} \, d\mu, \\ =& \sum_{k=0}^{\infty} \int_M \nabla_\xi \cdot Z \, d \circ \varphi^{-k} \, d\mu, \\ =& - \sum_{k=0}^{\infty} \int_M Z \cdot \left[\nabla_\xi \left(d \circ \varphi^{-k} \right) + \left(d \circ \varphi^{-k} \right) g \right] \, d\mu. \end{aligned} \tag{6.23}$$

Notice that now all integrands of the infinite series involve dot products between Z and two new vectors: the parametric gradient of $d \circ \varphi^{-k}$ and g re-scaled by the scalar $d \circ \varphi^{-k}$. The linear response related to the second vector is in fact a sum of time correlations between two observables. If the system is mixing and the property of exponential decorrelation holds, then

$$\left| \sum_{k=0}^{\infty} \int_M (d \circ \varphi^{-k}) Z \cdot g \, d\mu \right| \leq \exp(-C\, k), \tag{6.24}$$

where C denotes a positive constant whose value varies from system to system. From the definition of the unstable manifold, we use the fact that for any pair of points in M, x and y,

$$\text{dist}\left(\varphi^{-k}(x), \varphi^{-k}(y)\right) \to 0 \text{ as } k \to \infty. \tag{6.25}$$

This implies that norm of the parametric gradient of $d \circ \varphi^{-k}$ decays to zero with k. However, the rate of decay of individual components of that gradient will be different. The fastest decaying pairs are those that lie along the most expansive covariant direction. Thus, for a large k, the derivative

$$\partial_{\xi^i}\left(d \circ \varphi^{-k}(x)\right) = \lim_{\epsilon \to 0} \frac{d \circ \varphi^{-k}(x + \epsilon q^i) - d \circ \varphi^{-k}(x)}{\epsilon} \tag{6.26}$$

is smallest if $i = 1$.

Our observation of increasing components of the measure gradient combined with the analytical inspection of the unstable contribution, which was partially integrated yet again, leads to the following conclusion. Indeed, one could significantly decrease the magnitude of U by eliminating the following ingredients: 1) the slowest-decaying components of the parametric gradient of delayed d and, 2), components of g with a larger index. This could be achieved by choosing an *aligned* Z vector such that the linear response of the first term of the divergence expansion dominates over the remaining ones. This is true, for example, when 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.,

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

If J has such a Z vector, we say that J is an objective function *aligned* with q^1. In this special J, we could approximate U by keeping only the first term of $\nabla \cdot Z$. Thus,

$$U \approx -\sum_{k=0}^{\infty} \int_M Z^1 \left[\partial_{\xi^1} \left(d \circ \varphi^{-k}\right) + \left(d \circ \varphi^{-k}\right) g^1\right] d\mu := \sum_{k=0}^{\infty} U^k. \tag{6.27}$$

Note that by aligning the objective function, we are effectively integrating two scalar functions in each series term, $\partial_{\xi^1} \left(d \circ \varphi^{-k}\right)$ and $\left(d \circ \varphi^{-k}\right) g^1$. The second term is always proportional to g^1, which is typically a small number. In Lorenz 96 system, for example, $\|g^1\|_\infty \ll 1$ (see Figure 6-8). Also, notice that the leading term of the series, i.e., the one corresponding to the zero time-gap, can be upperbounded using the Cauchy-Schwartz inequality,

$$|U^1| \leq \|Z^1\|_2 \left(\|d\, g^1\|_2 \ + \|\partial_{q^1} d\|_2\right). \tag{6.28}$$

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 scattered positive Lyapunov spectrum as shown above. In the second term of Ineq. 6.28, d is differentiated in the most expansive direction q^1. It means that all components of the SRB density gradient 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 non-zero 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, by aligning J, the linear response related to the parametric gradient of the delayed function is reduced for the two following

reasons. The starting point of the series is significantly reduced and the rate of decay of subsequent terms is fastest possible.

In light of the specific behavior of the SRB density gradient and our analysis presented above, we shall numerically investigate the impact of the objective function J on the statistics and their change with respect to parameters. Here, we shall reconsider the two-dimensional version of the coupled sawtooth map introduced by Eq. 6.20. The purpose of this experiment is to visualize long-time averages computed at different parameter values for that 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 components 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, because $\partial_{q^2} Z^2 = 0$. 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 the only source of the divergence encoded in J 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 $z = x^1 - x^2$, $z = x^1$, and $z = x^1 + x^2$. The corresponding vector fields are such that

$$Z^1(z) = Z^2(z) = \int_C^z \exp(\sin(z'))\,\sin(z')\,dz', \tag{6.29}$$

with some arbitrary real-valued constant C. 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 sawtooth map. The statistics corresponding to these objective functions evaluated at a fine parametric grid are plotted in Figure 6-9. 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 Figure 6-9 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 could be true for any system's parameter with respect to which the sensitivity is computed. By applying integration by parts to the unstable contribution yet again, we show that the effect of larger components of measure gradient could be eliminated or reduced. Indeed, if there exists an aligned Z vector, we can still approximate the sensitivity accurately by removing larger components of g, cross-differentiating them in the direction of smallest slope, and removing the slow-decaying terms from the Ruelle's series. 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

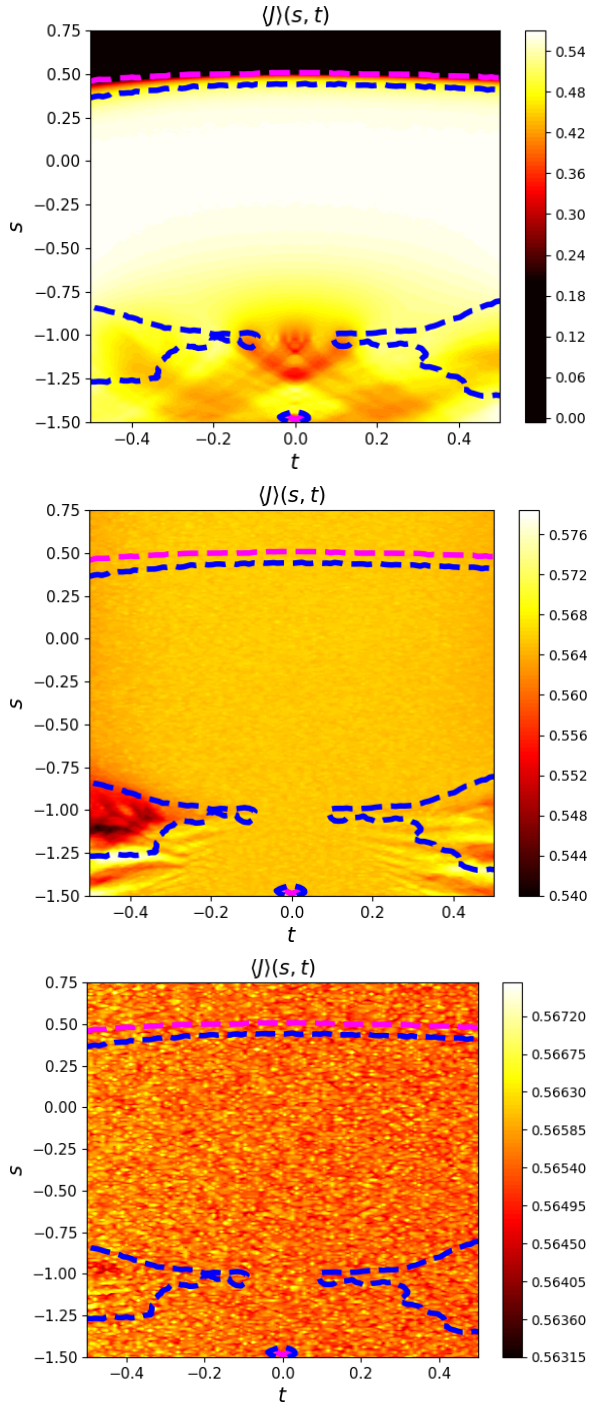

Figure 6-9: Long-time averages of the wave-like objective function $J = \exp(\sin(z))\sin(z)$, where $z = x^1 - x^2$ (upper plot), $z = x^1$ (middle plot) and $z = x^1 + x^2$ (lower plot). The time averages were computed for a uniform parametric grid consisting of 225 and 100 points along s and t, respectively. 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).

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. Consequently, we show that the stable part alone can approximate the total linear response sufficiently well.

6.4 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 Section 6.3, 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 Section 6.3, 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,

$$J = \frac{1}{b-a}\int_a^b \tilde{J}(x)\,dx \approx \frac{1}{n}\sum_{i=1}^{n}\tilde{J}(x^i) := \frac{1}{n}\sum_{i=1}^{n}\tilde{J}^i \tag{6.30}$$

where $\tilde{J} : \mathbb{R} \to \mathbb{R}$ 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

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

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

$$\langle J_w \rangle := \langle \sum_{i=1}^{n} w^i \, \tilde{J}^i \rangle = \sum_{i=1}^{n} \langle w^i \, \tilde{J}^i \rangle \overset{\text{indep.}}{=} \langle \tilde{J}^1 \rangle \langle \sum_{i=1}^{n} w^i \rangle = \langle J \rangle. \tag{6.31}$$

Eq. 6.31 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 $n \to \infty$.

We highlight yet another common property of larger physical systems. As reported by several publications (see [135] 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 [126], 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 steam,

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 [135], 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 [190], 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, $i = 2, ..., n$, 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" clusters 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. 6.31 holds.

6.4.1 Simplification of the S3 method

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

$$C = \int_M c^0 \, DJ \cdot f \, d\mu := \int_M (c^0 \, |f|) \, DJ \cdot q_f \, d\mu, \tag{6.32}$$

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. 6.6. 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, $i = 0, 1, ..., n$ 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 $n \to \infty$. 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 $n \to \infty$. 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 [15] originally proposed this approach for approximating LEs.

The total flop count of Algorithm 3 is proportional to m^2, which is a consequence of the recurrent QR factorization needed to find a basis of the center-unstable subspace every time step. This square dependence can be further alleviated if we introduce a skipping parameter that reduces the orthogonalization frequency of the basis matrix. Notice that if we set $v_{k+1} = r_{k+1}$, $Q_{k+1} = P_{k+1}$ and apply Lines 9-12 every k_{skip} time steps only, we do not violate the consistency of the linear response scheme

Algorithm 3: Reduced space-split sensitivity algorithm for higher-dimensional chaotic flows

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

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** $k = 0, ..., N-1$ **do** `// main time loop`
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**
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**

and effectively reduce the m-related factor of the total cost from m^2 to m^2/k_{skip}. While the aspect of cost reduction is straightforward, we will discuss the consistency of the modified scheme. By linearity and covariance property, if $k_{skip} > 1$, s can be decomposed into the original term and $\langle DJ \cdot \tilde{v}\rangle$, where $\tilde{v}$ denotes the center-unstable component of the inhomogeneous tangent. Using the assumption of small expectations of directional derivatives in any center-unstable direction, the extra term of sensitivity is negligible, analogously to the unstable and center contributions of the full S3 algorithm. To maintain the consistency of Benettin's algorithm, one must modify its output to $LEs := le/(N\,\Delta t\, k_{skip})$. This is because Line 6 is executed every time step and we are effectively adding accumulated logarithmic contribution of k_{skip} time steps. Notice that the R matrix of the modified scheme, which is now computed every k_{skip} time steps, is in fact a product of R matrices that would be generated by the standard reduced algorithm. In the case of upper-triangular matrices, the corresponding elements of the main diagonal are simply multiplied.

We also acknowledge that the optimal value of k_{skip} is a trade-off between the cost associated with the dimensionality of the unstable subspace and the ergodic-averaging

error. Indeed, by modifying k_{skip}, we also affect the latter. Recall that any standard tangent recursion $\tilde{v}$ grows in norm exponentially fast, i.e., $\|\tilde{v}_k\| \leq C \exp(\lambda_1 \Delta t\, k)$. Therefore, assuming statistical independence of spatial coordinates, the variance of $DJ{\cdot}\tilde{v}$ is upper-bounded by a function proportional to $\exp(2\,\lambda_1\,\Delta t\, k_{skip})$. Thus, finding optimal k_{skip} depends on the problem and requires some experimentation. Here, we set $k_{skip} = 1$. In the chapter on turbulence, we set $k_{skip} \gg 1$ due to large values of m and, consequently, long computation times.

6.4.2 Lorenz 96

In light of the above conclusions, we shall consider the Lorenz 96 model, which was proposed by E. Lorenz in [111] to study spatio-temporal dynamics of the atmosphere. Mathematically, this is an n-dimensional chaotic flow defined as follows,

$$\begin{aligned} \frac{dx^i}{dt} &= (x^{i+1} - x^{i-2})\, x^{i-1} - x^i + F, \quad i = 1, ..., n, \\ x^{i+n} &= x^i, \end{aligned} \tag{6.33}$$

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. Eq. 6.33 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., $x^i = F$, $i = 1, ..., n$ for all $t > t^* \approx 0$ [85]. We solve Eq. 6.33 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 Figure 6-10, 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

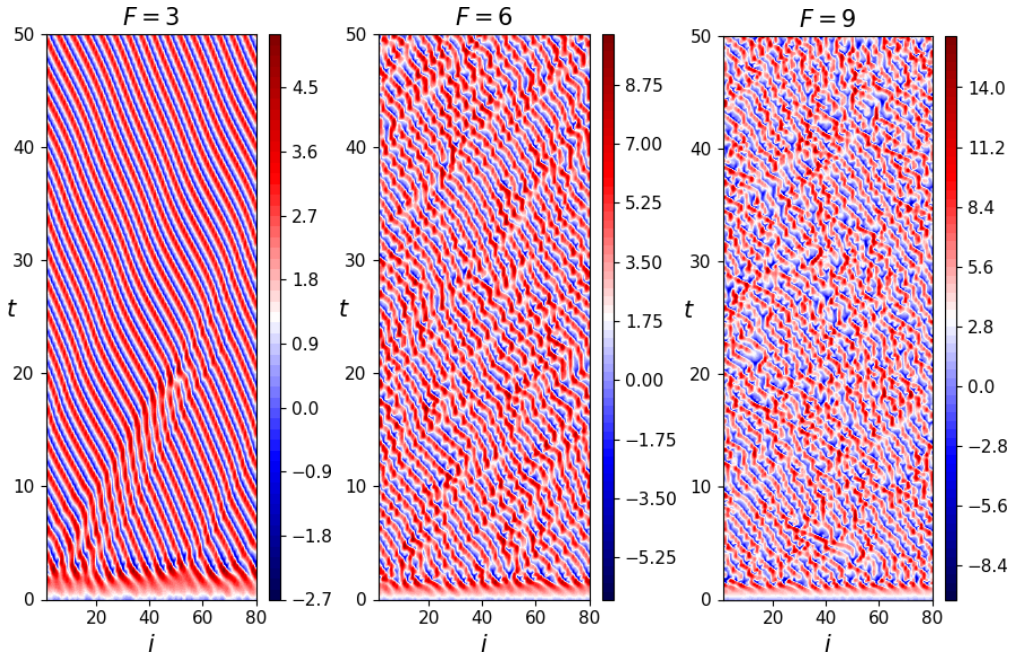

Figure 6-10: Solutions to the Lorenz 96 system (Eq. 6.33) for $n = 80$ stacked horizontally.

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 spatio-temporal 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 [176]. In Figure 6-11, 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 $\lambda_i(F)$, $i = 1, 2, ...$ and the x-axis. Indeed, the authors of [85] 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

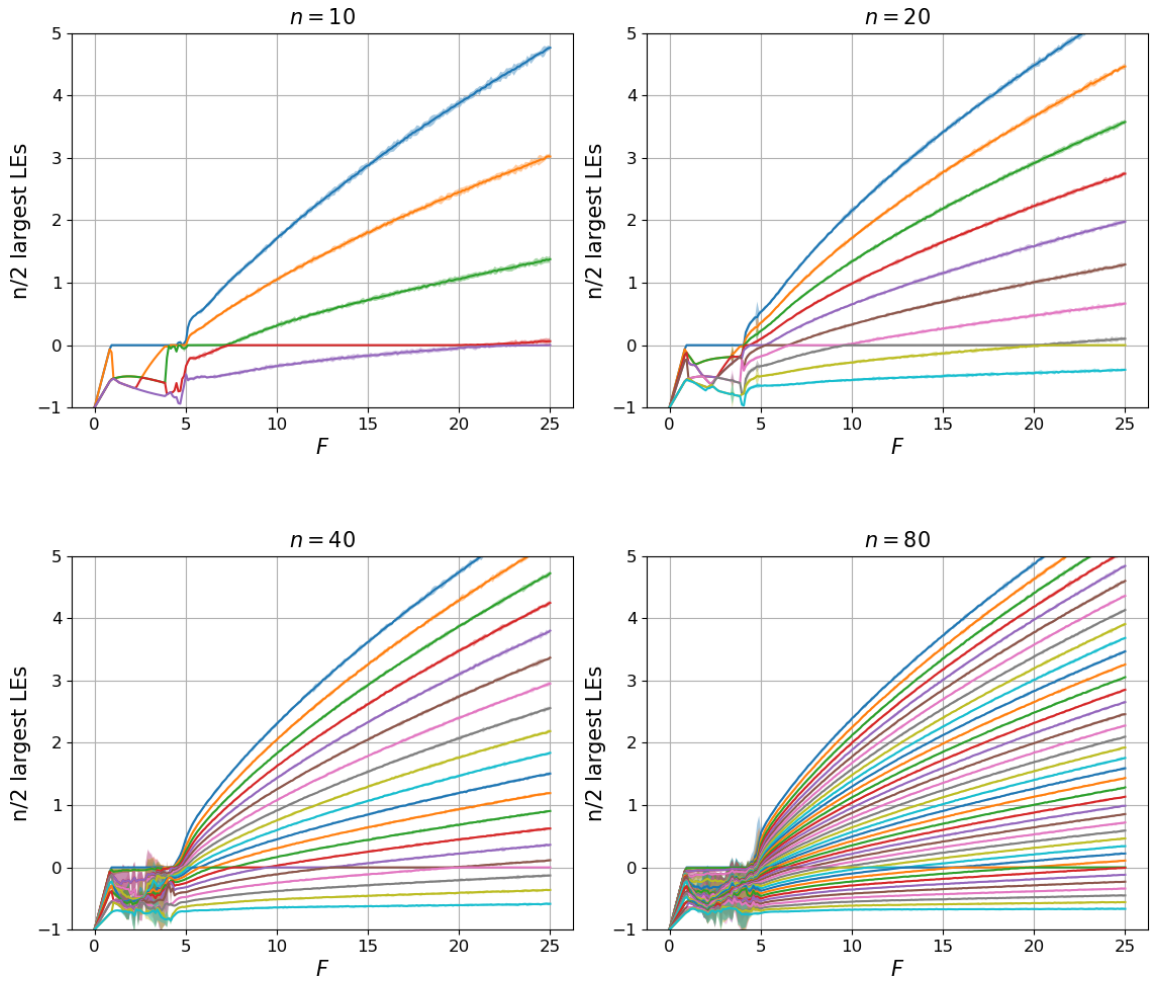

Figure 6-11: Larger half of the Lyapunov spectrum of Eq. 6.33. LEs were computed at 240 distinct values of F distributed uniformly between 0 and 25. For each value of F, we run 10 independent simulations over 5000 time units. The barely visible shaded area represents the 2-sigma range (95% confidence) of the 10-element data set at each value of F.

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. 6.30 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 Figure 6-12. 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 [85]. We shall restrict

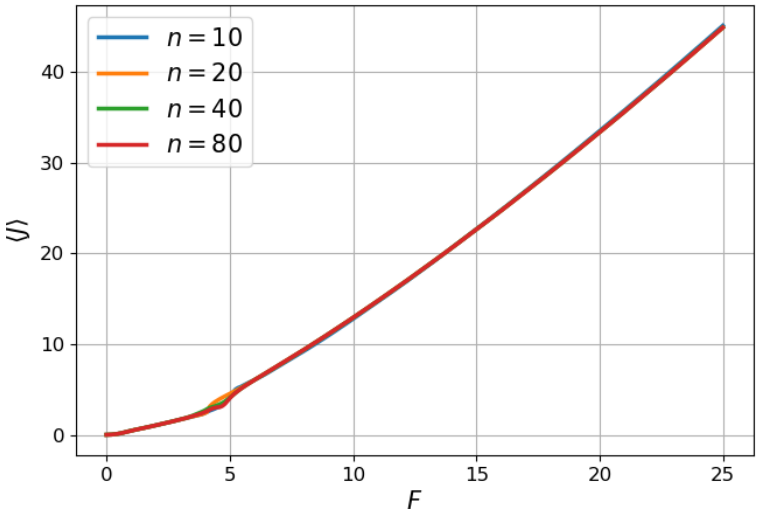

Figure 6-12: Long-time means of spatially-averaged kinetic energies of the Lorenz 96 system. The statistics were computed on a uniform grid of 240 values of $F \in [0, 25]$. For each value of F, the objective function was time-averaged over $5 \cdot 10^6$ time units.

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 3.

Figure 6-13 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 components 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 6-13 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 Section 6.3 suggests that in this case there might be no 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 Figure 6-14, we also plot the L^2 norms of the projection scalars c^i, $i = 1, ..., m_{ext} = m + 2$. 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.

6.4.3 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 spatio-temporal description of complex dynamics driven by instabilities far from an equilibrium. This equation was proposed decades ago to model wave propagation in

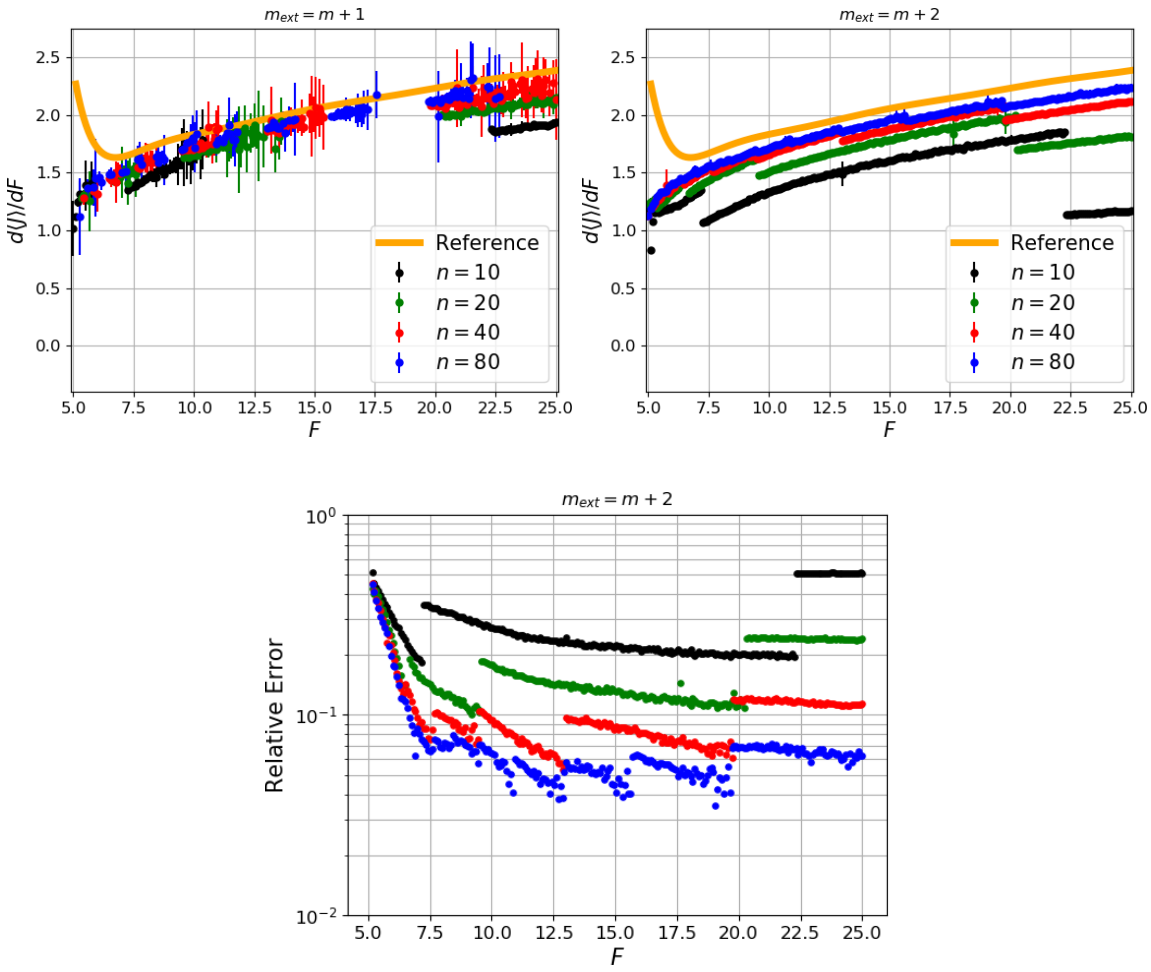

Figure 6-13: Linear response approximations of the Lorenz 96 model with respect to F computed using Algorithm 1. The upper left plot illustrates approximated sensitivities for $m_{ext} = m+1$, the upper right plot for $m_{ext} = m+2$, while the bottom plot depicts the mean relative error of the $m_{ext} = m + 2$ 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 $\langle J\rangle$.

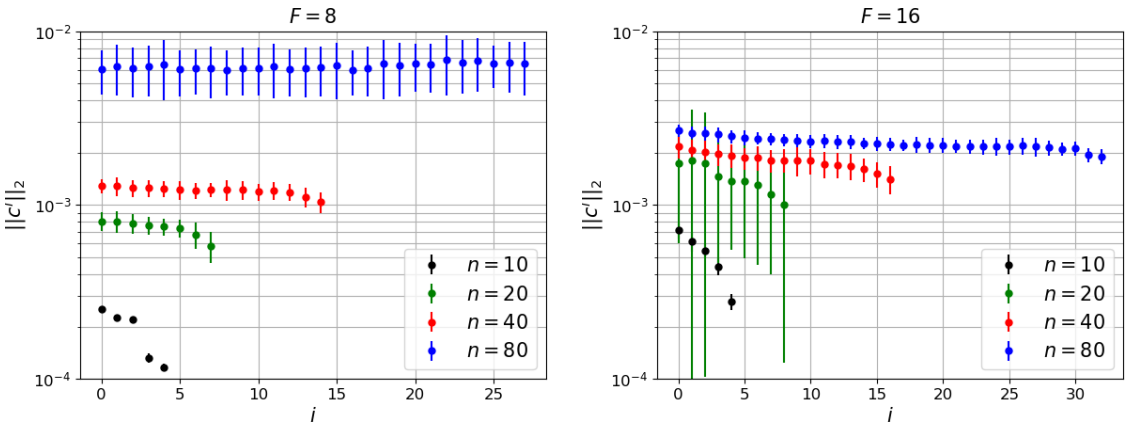

Figure 6-14: L^2 norms of c^i, $i = 1, ..., m_{ext} = m + 2$, which were computed as by-products of Algorithm 1. All simulation parameters are the same as those reported in the caption of Figure 6-13.

reaction-diffusion systems [91] and hydrodynamic instabilities of laminar flames [156]. 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 [21], has the following form,

$$\begin{aligned} &\frac{\partial u}{\partial t} = -(u + c)\frac{\partial u}{\partial x} - \frac{\partial^2 u}{\partial x^2} - \frac{\partial^4 u}{\partial x^4}, \\ &u(0, t) = u(L, t) = 0, \\ &\frac{\partial u}{\partial x}(0, t) = \frac{\partial u}{\partial x}(L, t) = 0, \end{aligned} \tag{6.34}$$

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 [21]. 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 6.6.2, we discuss how the linear response algorithm could be integrated with implicit schemes.

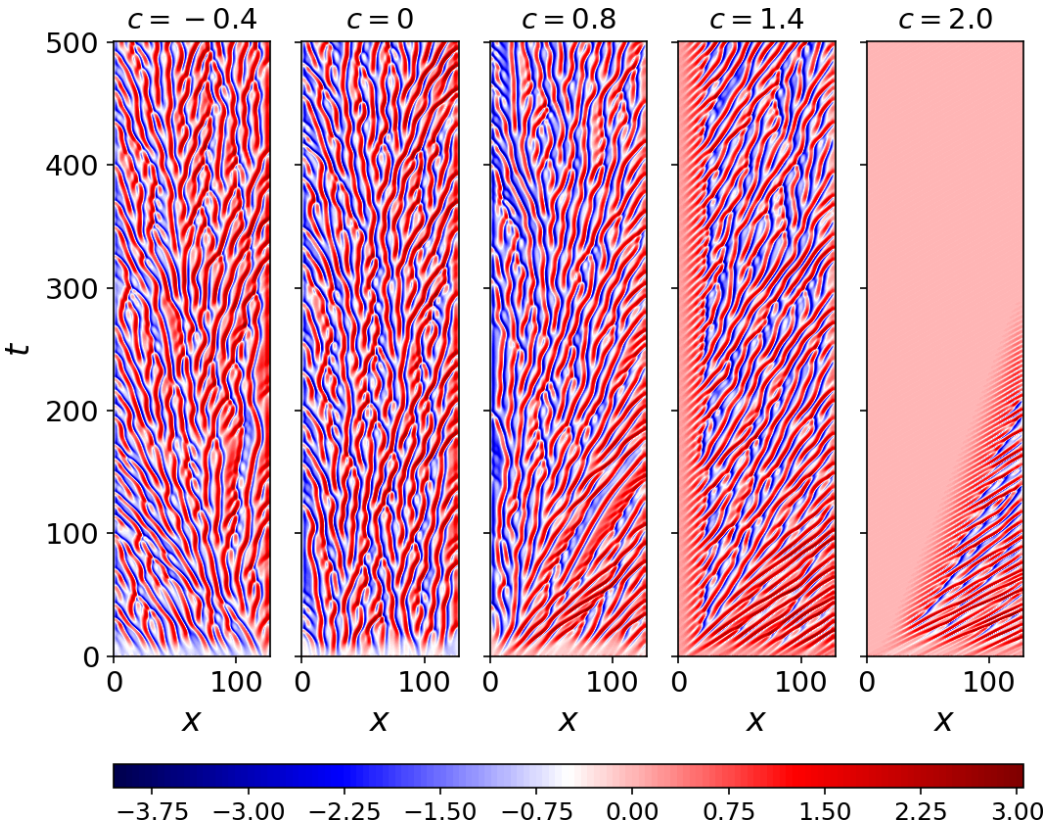

Figure 6-15: Solutions to the Kuramoto-Sivashinsky equation (Eq. 6.34) for different advection intensities.

Figure 6-15 illustrates solutions to the KS equation, $u(x,t)$, for different values of c. In the spatio-temporal 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 spatio-temporal 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 6-16 depicts the 18 largest Lyapunov exponents of the KS equations for

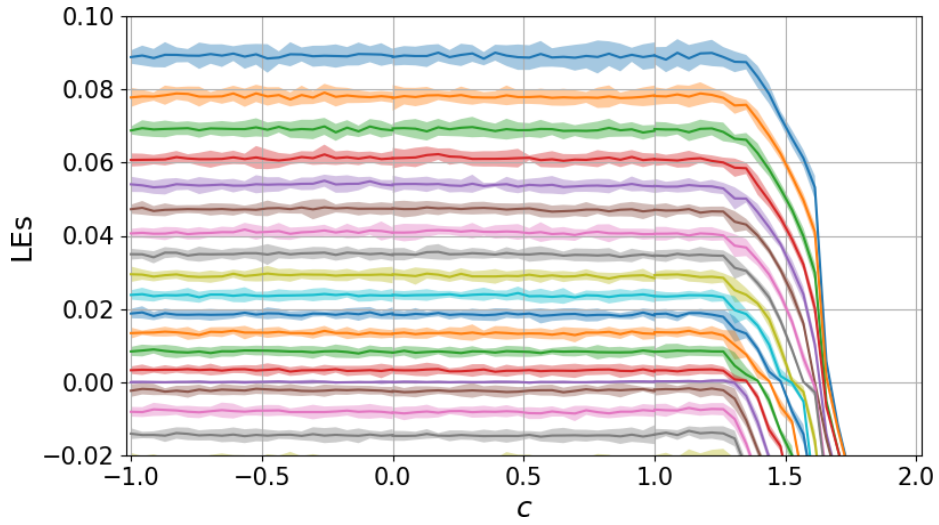

Figure 6-16: 18 largest Lyapunov exponents of the KS equation. The spectrum was computed at the uniform grid between $c = -1$ and $c = 2$. For each value of c, 10 independent simulations were run. The sought-after quantities were obtained through ergodic-averaging over $12,000$ time units per simulation. The solid lines represent the mean values obtained in 10 simulations, while the shaded area represents the 2-sigma range.

$c \in [-1, 2]$. 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 Figure 6-15.

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 Figure 6-17. 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 spatio-temporal domain. The tangent vectors corresponding to moderate indices are placed in the bottom row of Figure 6-17. 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 [161, 15], 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 [89, 92, 169].

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 Figure 6-18 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 Figure 6-15: turbulent ($c \leq 1.25$), transitional

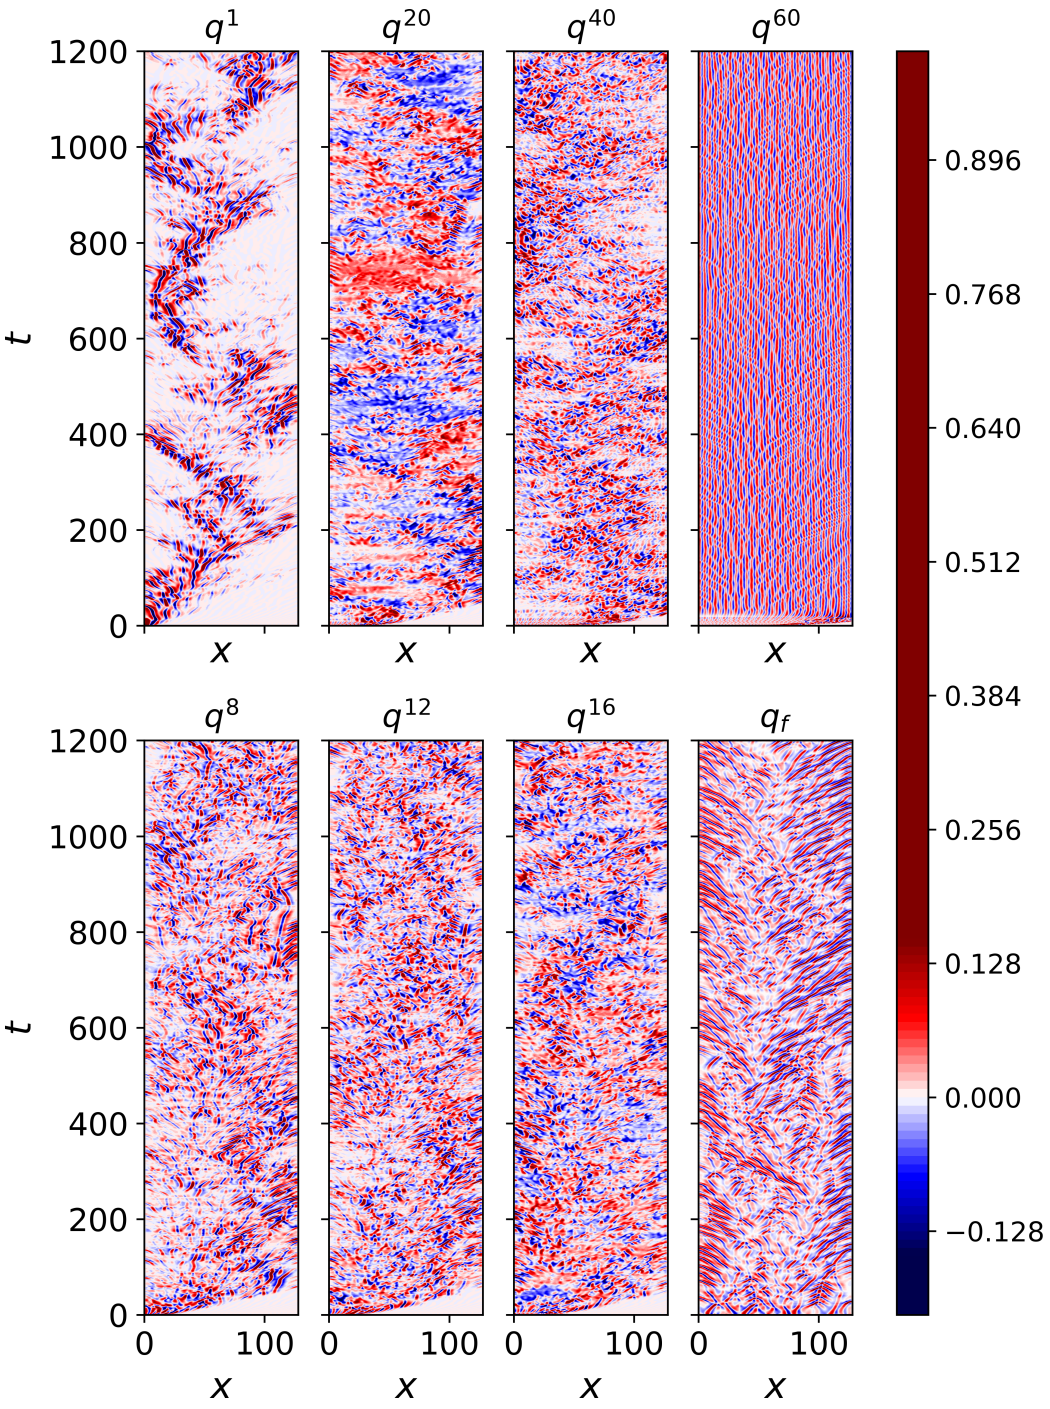

Figure 6-17: Orthonormal Lyapunov vectors q^i of the KS system (Eq. 6.34) without the extra advection term ($c = 0$). The vector q_f represents the normalized time derivative of $u(x,t)$. The colorbar has linearly been re-scaled between -0.15 and 0.15 keeping the same color for all values from beyond this interval.

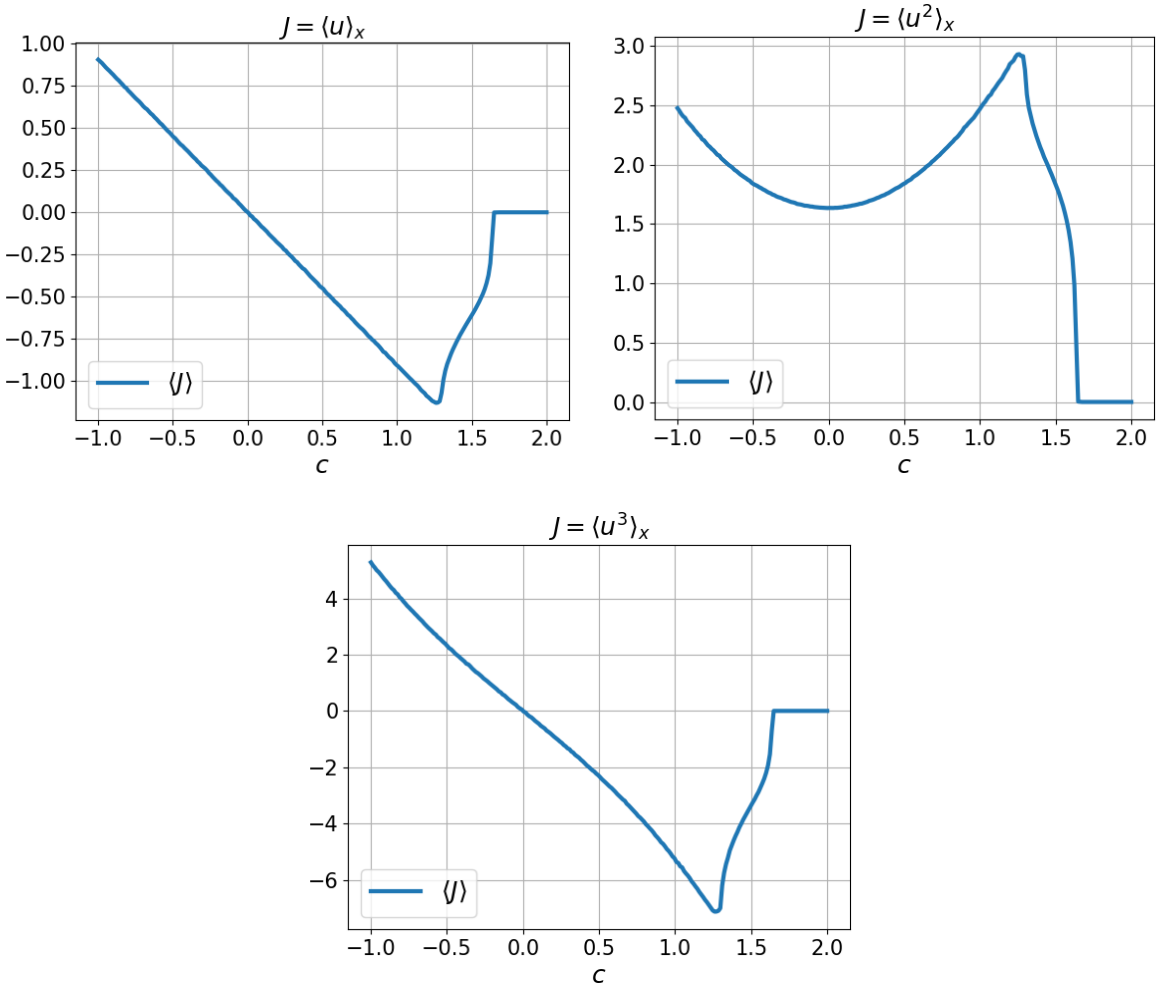

Figure 6-18: Long-time averages $\langle J \rangle$ computed on a uniform 240-point grid of $c \in [-1, 2]$. The operator $\langle \cdot \rangle_x$ indicates the spatial average. For each value of c, we run an ergodic-averaging simulation over 600,000 time units.

$(1.25 \leq c \leq 1.7)$, and advection-dominated $(c \geq 1.7)$ regime.

We apply our reduced linear response algorithm (Algorithm 1) to approximate sensitivities for these three objective functions. Analogously to the previous plots, we compare our approximations against the finite-difference reference solutions. Figure 6-19 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 Figure 6-13, we observed that the approximation error decreases as $n \to \infty$. Indeed, the dimension of the 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 [129], 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 broad family of time-delayed dynamical systems. The spatio-temporal structure of the laser dynamics with delayed feedback presented in [5, 62] clearly features a statistically homogeneous behavior. The user would need to represent such a system using an appropriate diffeomorphic map $\varphi : M \to M$ and compute relevant phase-space and parametric derivatives, following the recipe described in this chapter. 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 [5]. In an analogous way, one can easily derive φ for any non-autonomous system.

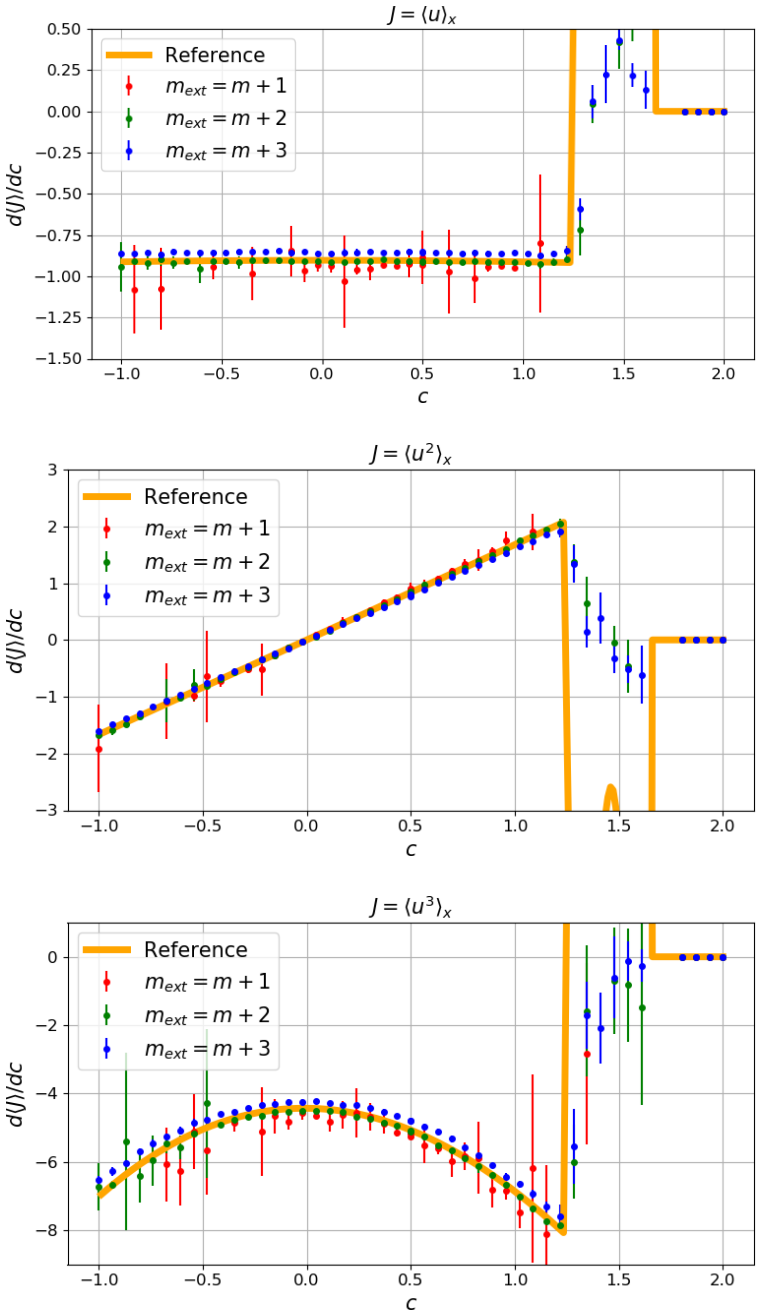

Figure 6-19: Linear response computed for the same objective functions as those presented in Figure 6-18 using Algorithm 1. For each value of c, we run 10 independent simulations over 3,000 time units each. Bullets and vertical lines represent the mean and standard deviation, respectively. The results with a large standard deviation were removed from the plot. The reference line was computed through central finite-differencing of polynomial fits.

6.5 Summary

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.

6.6 Appendices

This section includes supplementary material supporting certain arguments presented in this chapter.

6.6.1 Full space-split algorithm – description, pseudocode and complexity analysis

The purpose of this section is to extend the discrete version of S3 derived in Chapter 5 to continuous chaos and present the structure of the full linear response algorithm. We rely on the three-term splitting defined by Eq. 6.3. 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, $i = 0, 1, ..., m$, 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 Section 6.2. Another consequence of the three-term splitting is the emergence of the neutral contribution of the linear response. Fortunately, as shown in Eq. 6.10-6.14, 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. Eq. 6.16 indicates that we need c^i, $i = 0, 1, ..., m$, 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 [162] 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. 6.16) involves only self-derivatives of c^i, i.e., $b^{i,j}$ 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 [162]. 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: $x_k(\xi) : [0,1]^m \to U_k \subset M$ 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 $\nabla_{\xi_k} x_k = [\partial_{\xi_k^1} x_k, ..., \partial_{\xi_k^m} x_k]$. A more rigorous description and analysis of this coordinate transformation can be found in [162].

To obtain $b^{i,j}$, $i = 0, 1, ..., m$, $j = 1, ..., m$, we simply differentiate Eq. 6.5, Eq. 6.7 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. 6.5 with respect to ξ_{k+1}^j yields

$$\begin{aligned} &\partial_{\xi_{k+1}^j} v_{k+1} := w_{k+1}^j = \\ &\partial_{\xi_{k+1}^j} r_{k+1} - \sum_{l=1}^{m} b_{k+1}^{l,j}\, q_{k+1}^l + c_{k+1}^l\, p^{l,j} - b_{k+1}^{0,j}\, f_{k+1} - c_{k+1}^0\, Df_{k+1}\, q_{k+1}^j, \end{aligned} \tag{6.35}$$

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. 6.7, i.e., constraint enforcing $v \cdot q = 0$, with respect

to ξ_{k+1}^j gives

$$\begin{aligned} b_{k+1}^{i,j} &= p_{k+1}^{i,j} \cdot \left(r_{k+1} - c_{k+1}^0 \, f_{k+1}\right) + q_{k+1}^i \cdot \partial_{\xi_{k+1}^j} r_{k+1} \\ &\quad - b_{k+1}^{0,j} q_{k+1}^i \cdot f_{k+1} - c_{k+1}^0 \, q_{k+1}^i \cdot Df_{k+1} \, q_{k+1}^j. \end{aligned} \tag{6.36}$$

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. 6.35 to obtain

$$\begin{aligned} &- v_{k+1} \cdot Df_{k+1} \, q_{k+1}^j = w_{k+1}^j \cdot f_{k+1} = \\ &\partial_{\xi_{k+1}^j} r_{k+1} \cdot f_{k+1} - \sum_{l=1}^{m} b_{k+1}^{l,j} \, q_{k+1}^l \cdot f_{k+1} + c_{k+1}^l \, p^{l,j} \cdot f_{k+1} \\ &- b_{k+1}^{0,j} \, f_{k+1} \cdot f_{k+1} - c_{k+1}^0 \, Df_{k+1} \, q_{k+1}^j \cdot f_{k+1}. \end{aligned} \tag{6.37}$$

Finally, by combining Eq. 6.36–6.37, we derive the following linear system for $b^{i,j}$, $i = 0, 1, ..., m$, $j = 1, ..., m$,

$$\begin{aligned} &(f_{k+1} \cdot f_{k+1}) b_{k+1}^{0,j} + \sum_{l=1}^{m} (q_{k+1}^l \cdot f_{k+1}) \, b_{k+1}^{l,j} = d_{k+1}^{0,j}, \\ &(q_{k+1}^i \cdot f_{k+1}) b_{k+1}^{0,j} + b_{k+1}^{i,j} = d_{k+1}^{i,j}, \;\; i, j = 1, ..., m, \end{aligned} \tag{6.38}$$

where

$$\begin{aligned} d_{k+1}^{0,j} &:= v_{k+1} \cdot Df_{k+1} \, q_{k+1}^j + \partial_{\xi_{k+1}^j} r_{k+1} \cdot f_{k+1} \\ &- \sum_{l=1}^{m} c_{k+1}^l \, p_{k+1}^{l,j} \cdot f_{k+1} - c_{k+1}^0 \, Df_{k+1} \, q_{k+1}^j \cdot f_{k+1}, \\ d_{k+1}^{i,j} &:= p_{k+1}^{i,j} \cdot \left(r_{k+1} - c_{k+1}^0 \, f_{k+1}\right) + q_{k+1}^i \cdot \partial_{\xi_{k+1}^j} r_{k+1} \\ &- c_{k+1}^0 \, q_{k+1}^i \cdot Df_{k+1} \, q_{k+1}^j, \;\; i, j = 1, ..., m. \end{aligned} \tag{6.39}$$

The Schur complement of System 6.38–6.39 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,

$$\begin{aligned} b_{k+1}^{i,j} &= (S_{k+1}^{-1})^{i:} \cdot d_{k+1}^{1:m,j} - \sum_{l=1}^{m} (S_{k+1}^{-1})^{il} \frac{q_{k+1}^{l} \cdot f_{k+1}}{f_{k+1} \cdot f_{k+1}} d_{k+1}^{0,j} \\ &= (S_{k+1}^{-1})^{i:} \cdot \left(d_{k+1}^{1:m,j} - \frac{d_{k+1}^{0,j}}{f_{k+1} \cdot f_{k+1}} Q_{k+1}^{T} f_{k+1} \right), \; i,j = 1, ..., m, \end{aligned} \tag{6.40}$$

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 $i = 1, ..., m$ and a fixed j. Once $b^{i,j}$ for all $i, j = 1, ..., m$ is computed, $b^{0,j}$ and w^{j}, $j = 1, ..., m$ can be evaluated directly using Eq. 6.35 and Eq. 6.38.

Based on Eq. 6.35–6.40, 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 [161]. Using the definition of r_{k+1} and all underlying quantities, we apply the chain rule to expand $\partial_{\xi_{k+1}^{j}} r_{k+1}$,

$$\partial_{\xi_{k}^{j}} r_{k+1} = D^2\varphi(v_k, q_k^j) + D\varphi_k \, w_k^j + D\partial_s\varphi_k \, q_k^j, \tag{6.41}$$

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. 6.41 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 [162]. Based on the above derivations, Section 6.2 and [161], 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.

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 6.6.2, 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

$$\mathcal{O}\left(N\,n^3\,(m^2 + n_s\,m)\right). \tag{6.42}$$

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. There-

Algorithm 4: Space-split sensitivity algorithm for hyperbolic flows

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

Randomly generate: x_0, v_0, Q_0, $a_0^{i,j}$, w_0^i for all $i,j=1,...,m$;
Set $s=c=u=0$;
for $k=0,...,N-1$ **do** `// main time loop`
  if $k \geq T$ **then**
    $s := s + DJ_k \cdot v_k$;
    $u := u - J_k\,(u_k + u_{k-1} + ... + u_{k-K+1})$;
    $c := c + DJ_k \cdot f_k\,(c_k^0 + c_{k-1}^0 + ... + c_{k-K+1}^0)$;
  end
  Update stable (s), neutral (c) and unstable (u) contributions [28, 161].
  $P_{k+1} = D\varphi_k\, Q_k$;
  QR-factorize P_{k+1}: $Q_{k+1}\,R_{k+1} = P_{k+1}$;
  Find the inverse of R_{k+1};
  Update basis vectors Q and transformation Jacobian R. See [51] for derivation and convergence analysis of Lyapunov basis.
  for $i=1,...,m,\ j=1,...,i$ **do** $\tilde{a}_{k+1}^{i,j} = D^2\varphi_k(q_k^i, q_k^j) + D\varphi_k\, a_k^{i,j}$;
  for $i=1,...,m,\ j=1,...,i$ **do** $a_{k+1}^{i,j} = \tilde{a}_{k+1}^{p,q}\,(R^{-1})_{k+1}^{pi}\,(R^{-1})_{k+1}^{qj}$;
  for $i=1,...,m$ **do**
    for $p,q=1,...,m$ **do**
      $$(\partial_{\xi_{k+1}^i} R_{k+1})^{pq} = \begin{cases} q_{k+1}^p \cdot a_{k+1}^{p,i}, & \text{if } p=q \\ q_{k+1}^p \cdot a_{k+1}^{q,i} + q_{k+1}^q \cdot a_{k+1}^{p,i}, & \text{if } p<q \\ 0, & \text{otherwise} \end{cases};$$
    end
    $g_{k+1}^i = -\mathrm{tr}(\partial_{\xi_{k+1}^i} R_{k+1})$;
  end
  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 [162] for original derivation.
  for $i,j=1,...,m$ **do** $p_{k+1}^{i,j} = a_{k+1}^{i,j} - q_{k+1}^l(\partial_{\xi_{k+1}^j} R_{k+1})^{li}$;
  Compute derivatives of Lyapunov vectors [161].
  $S_{k+1} = I - Q_{k+1}^T f_{k+1}(Q_{k+1}^T f_{k+1})^T / f_{k+1}\cdot f_{k+1}$;
  Find the inverse of S_{k+1};
  $r_{k+1} = D\varphi_k\, v_k + \chi_{k+1}$;
  $z_{k+1} = Q_{k+1}^T\,(r_{k+1} - (f_{k+1}\cdot r_{k+1})/(f_{k+1}\cdot f_{k+1}) f_{k+1})$;
  for $i=1,...,m$ **do** $c_{k+1}^i = (S_{k+1}^{-1})^{ij}\, z_{k+1}^j$;
  $c_{k+1}^0 = f_{k+1}\cdot(r_{k+1} - c_{k+1}^i\, q_{k+1}^i)/f_{k+1}\cdot f_{k+1}$;
  $v_{k+1} = r_{k+1} - c_{k+1}^i\, q_{k+1}^i - c_{k+1}^0\, f_{k+1}$;
  for $i=1,...,m$ **do** $\partial_{\xi_k^i} r_{k+1} = D^2\varphi_k(v_k, q_k^i) + D\varphi_k\, w_k^i + D\partial_s\varphi_k\, q_k^i$;
  $\nabla_{\xi_{k+1}} r_{k+1} = \nabla_{\xi_k} r_{k+1}\, R_{k+1}^{-1}$;

  for $i=1,...,m$ **do**
    $d_{k+1}^{0,i} = v_{k+1}\cdot Df_{k+1}\, q_{k+1}^i + \partial_{\xi_{k+1}^i} r_{k+1}\cdot f_{k+1} - c_{k+1}^l\, p_{k+1}^{l,i}\cdot f_{k+1} -$
    $c_{k+1}^0\, Df_{k+1}\, q_{k+1}^i \cdot f_{k+1}$;
  end
  for $i,j=1,...,m$ **do**
    $d_{k+1}^{i,j} = p_{k+1}^{i,j}\cdot(r_{k+1} - c_{k+1}^0\, f_{k+1}) + q_{k+1}^i\cdot\partial_{\xi_{k+1}^j} r_{k+1} - c_{k+1}^0\, q_{k+1}^i\cdot Df_{k+1}\, q_{k+1}^j$;
  for $i,j=1,...,m$ **do** $b_{k+1}^{i,j} = (S_{k+1}^{-1})^{i:}\cdot(d_{k+1}^{1:m,j} - d_{k+1}^{0,j}/(f_{k+1}\cdot f_{k+1}) Q_{k+1}^T f_{k+1})$;
  for $i=1,...,m$ **do** $b_{k+1}^{0,i} = 1/(f_{k+1}\cdot f_{k+1})\,(d_{k+1}^{0,i} - (q_{k+1}^l\cdot f_{k+1}) b_{k+1}^{l,i})$;
  for $i=1,...,m$ **do**
    $w_{k+1}^i = \partial_{\xi_{k+1}^i} r_k - b_{k+1}^{l,i}\, q_{k+1}^l - c_{k+1}^l\, p_{k+1}^{l,i} - b_{k+1}^{0,i}\, f_{k+1} - c_{k+1}^0\, Df_{k+1} q_{k+1}^i$;
  Modification of the general discrete S3 from [161]. This part computes c, b, w, according to the derivation presented in this section and Section 2.1.
  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;
  Advance the iteration: $x_{k+1} = \varphi(x_k)$;
end

fore, the leading term of the flop count dramatically decreases to

$$\mathcal{O}\left(N\, n\left(m^3 + n_s\, m^2 + n_s\, n_J\right)\right). \tag{6.43}$$

6.6.2 Handling non-trivial time integrators 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 work, 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. 1.4. 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

$$\varphi(x_k) = x_k + \Delta t\, f\left(x_k + \frac{\Delta t}{2}\, f(x_k)\right) := x_k + \Delta t\, f(x_p) = x_k + \Delta t\, f_p, \tag{6.44}$$

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

$$D\varphi_k\, v = v + \Delta t\, Df_p\, v + \frac{\Delta t^2}{2}\, Df_p\, Df_k\, v, \tag{6.45}$$

with $Df_k = Df(x_k)$ and $Df_p = Df(x_p)$, in compliance with our notation convention. Differentiating Eq. 6.45 once more and contracting it against yet another vector

$a \in \mathbb{R}^n$, we obtain the following relation,

$$\begin{aligned} D^2\varphi_k(v,a) =& \Delta t\, D^2 f_p \left(v + \frac{\Delta t}{2}\, Df_k\, v, a \right) + \\ & \frac{\Delta t^2}{2}\, D^2 f_p \left(v + \frac{\Delta t}{2}\, Df_k\, v, Df_k\, a \right) + \\ & \frac{\Delta t^2}{2}\, Df_p\, D^2 f_k(v,a). \end{aligned} \tag{6.46}$$

Recall that $D^2\varphi_k(v,a) \in \mathbb{R}^n$. Assuming f also depends on a scalar parameter s, the parametric derivative of Eq. 6.44 expands as follows,

$$\partial_s \varphi_k = \chi_{k+1} = \Delta t\, \partial_s f_p + \frac{\Delta t^2}{2} Df_p\, \partial_s f_k, \tag{6.47}$$

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. 6.47,

$$\begin{aligned} D\partial_s\varphi_k\, v =& \Delta t\, D\partial_s f_p \left(v + \frac{\Delta t}{2}\, Df_k\, v \right) + \\ & \frac{\Delta t^2}{2} D^2 f_p \left(v + \frac{\Delta t}{2}\, Df_k\, v, \partial_s f_k \right) + \\ & \frac{\Delta t^2}{2}\, Df_p\, D\partial_s f\, v. \end{aligned} \tag{6.48}$$

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-vectors 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. 6.44–6.48 are computed. Let us consider a generic implicit scheme,

$$h(x_k, x_{k+1}) = 0, \tag{6.49}$$

where $x_{k+1} = \varphi(x_k)$. Assuming x_k is known, the n-dimensional nonlinear system defined by Eq. 6.49 is typically solved for x_{k+1} using a standard solver such as the Newton-Raphson method. Differentiating Eq. 6.49 with respect to x_k and multiplying both sides by a vector v, we obtain the following system,

$$\frac{\partial h}{\partial x_{k+1}} D\varphi_k \, v = -\frac{\partial h}{\partial x_k} \, v, \tag{6.50}$$

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. 6.50 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. 6.50, apply the chain rule as presented above, and formulate analogous linear systems.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 7

Tangent dynamics and sensitivity analysis of turbulent flows

Although there has recently been huge progress in the sensitivity analysis of chaos, only a small handful of new methods has been applied to three-dimensional PDEs. In the development process, the accuracy of the linear response approximation is only one among several important aspects that needs to be carefully studied. Other include Lyapunov spectra, dynamics of Lyapunov vectors and impact of the mesh resolution. They are helpful in understanding the underlying physics and calibrating the existing methods. Here, we shall exclusively focus on turbulent channel flows. The purpose of this chapter is twofold. First, we analyze aspects of the underlying tangent dynamics that play a big role in the computation of sensitivities. Second, we apply the reduced algorithm for the sensitivity analysis proposed in Chapter 6 and assess its performance. Numerical results presented in this chapter is a starting point in the development of high-fidelity linear response machinery for real-world turbulent flows.

7.1 Compressible isothermal channel flow: problem description

Let us consider the compressible isothermal Navier-Stokes system defined by the following two coupled PDEs,

$$\begin{aligned} &\frac{\partial \rho}{\partial t} + \frac{\partial (\rho\, u_j)}{\partial x_j} = 0, \\ &\frac{\partial (\rho u_i)}{\partial t} + \frac{\partial (\rho\, u_i\, u_j)}{\partial x_j} + \frac{\partial p}{\partial x_i} - \frac{\partial}{\partial x_k}\left(\nu\, \rho \left(\frac{\partial u_i}{\partial x_k} + \frac{\partial u_k}{\partial x_i}\right)\right) = 0, \end{aligned} \tag{7.1}$$

which represent the mass and momentum conservation laws, respectively. This system is closed with the algebraic equation of state,

$$p = \rho\, c^2, \tag{7.2}$$

where c is the (uniform in space) speed of sound, u_i, $i = 1, 2, 3$ are the components of the velocity vector, while ρ denotes the density of the fluid. The only input/control parameter is the kinematic viscosity denoted by ν, which is a positive scalar. Repeated subscripts imply summation over all three spatial dimensions per Einstein's convention. In this chapter, we will be analyzing an *approximate* version of System 7.1–7.2. In particular, the original diffusive term (i.e., the term proportional to ν) is replaced with

$$-\nu\, \rho\, \frac{\partial}{\partial x_k}\left(\frac{\partial u_i}{\partial x_k}\right) - \nu_{bulk}\, \rho\, \frac{\partial}{\partial x_i}\left(\frac{\partial u_k}{\partial x_k}\right), \tag{7.3}$$

where $i = 1, 2, 3$. Thus, some parts of the gradient of the diffusivity are neglected here. In addition, we add a bulk diffusivity proportional to a new input parameter ν_{bulk} to eliminate acoustic waves from the flow. None of these changes affects the turbulent behavior of the system. In all our numerical experiments, we set ν_{bulk} to 0.9.

This modified Navier-Stokes system is defined in a box $[0, L_1] \times [-L_2/2, L_2/2] \times [0, L_3]$, where 1, 2 and 3 are the labels of Cartesian coordinates that represent the streamwise, wall-normal and spanwise directions, respectively. We will alternately

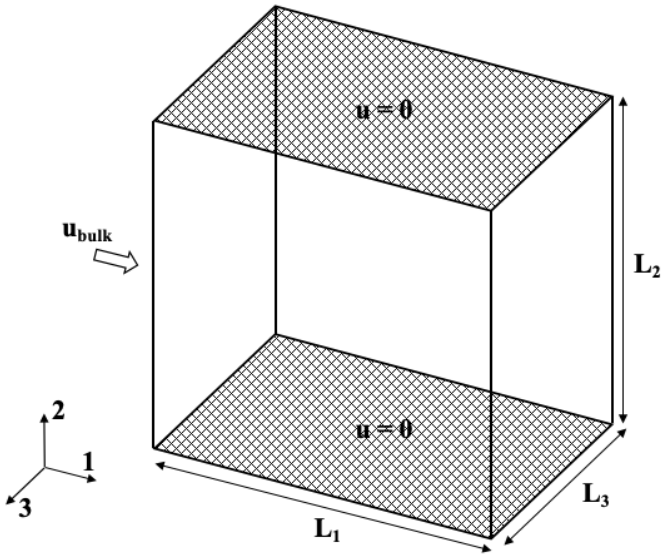

Figure 7-1: Sketch of the computational domain used for the channel flow simulation.

use x, y, z to denote the corresponding spatial directions. The domain parameters are fixed throughout this chapter and they equal: $L_1 = 2$, $L_2 = 2\delta = 2$, $L_3 = 1$. See Figure 7-1 for a sketch of the computational domain considered in this study.

The boundary conditions are periodic in x and z for all quantities, and $u_i = 0$, $i = 1, 2, 3$, $\partial\rho/\partial y = 0$ at $y \pm \delta$. The initial condition is zero except the streamwise component of the velocity, which consists of two ingredients: a quadratic-function laminar profile and random solenoidal perturbation.

The continuous governing equations are discretized in space using second-order finite differences. The grid is uniform in both x and z directions, and non-uniform in the y direction. The grid spacing along the periodic directions is L_1/n_1 and L_3/n_3, where n_1 and n_3 respectively denote the number of distinct points in the x and z directions. In the wall-normal direction, we generate n_2 points with refinement close to the wall. The semi-discrete system of ODEs is integrated in time using the explicit fourth-order Runge-Kutta scheme.

Note that ν is constant throughout the domain under the assumed isothermal condition. Consequently, the product of the temperature and the gas constant is constant as well. We set its value to 100. The bulk velocity vector is aligned with the first coordinate and equals 1, which means that the Mach number of the flow equals

0.1, while the bulk Reynolds number

$$Re_{bulk} = \frac{u_{bulk}\,\delta}{\nu} = \frac{1}{\nu}. \tag{7.4}$$

Note also that although the density varies in space, the small value of Mach number implies the compressibility effects are expected to have low impact on the flow physics.

The reader is referred to [132] and the `pyChan` section of [133] for further details on the construction of the initial condition, grid distribution in the wall-normal direction and time step adjustment scheme.

We run a test simulation using the numerical framework described above and illustrate some physical aspects of the flow. In this test case, we set $\nu = 3 \cdot 10^{-3}$ and choose a mesh with $(n_1, n_2, n_3) = (25, 100, 25)$ such that the grid resolution in the periodic directions equals $\Delta x^+ = 17.2$, $\Delta z^+ = 8.6$, while the wall-normal spacing is $\Delta y^+ = 1.1$ and $\Delta y^+ = 7.6$ at the wall and channel center, respectively. Any wall-normal distance l can be expressed using *wall units*, represented by the superscript "+", through the following formula: $l^+ = u_\tau\, l/\nu$, where u_τ denotes the friction velocity. In this particular simulation, $u_\tau \approx 0.0619$, which implies that the corresponding friction Reynolds number, defined as

$$Re_\tau = \frac{u_\tau\,\delta}{\nu}, \tag{7.5}$$

is approximately 206. Figure 7-2 illustrates a select time snapshot of the streamwise velocity reflecting typical structure of the flow, while Figure 7-3 plots some time-averaged quantities obtained by running the test simulation over $t^+ = u_\tau t/\delta \approx 124$ time units.

The wall-normal profiles of u_2 and u_3 are almost zero due to the specific boundary conditions, while the one of density mildly oscillates close to the value of one confirming the almost incompressible behavior of the flow. We also observe that the mean profile of the streamwise component is consistent with the theory (see Section 7 of [140]). In particular, the mean velocity is clearly linear with respect to y^+ in the viscous sublayer, i.e., when $y^+ < 5 - 7$, and features logarithmic behavior in the

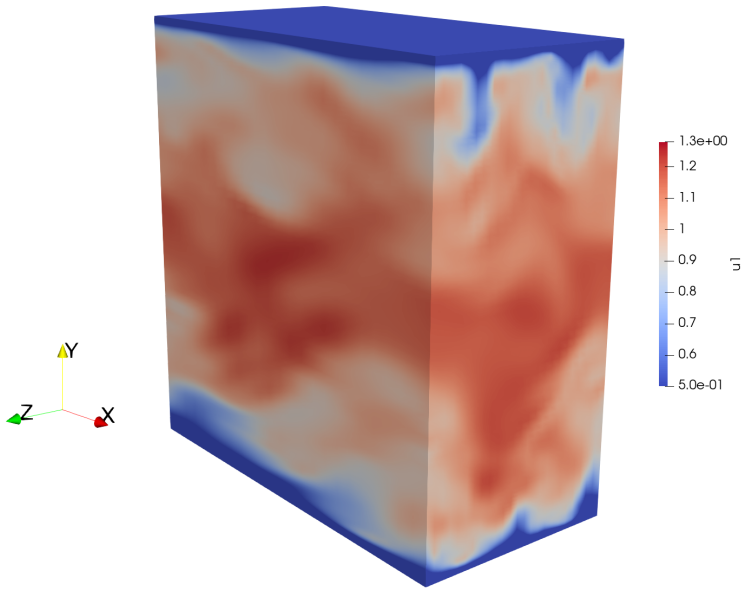

Figure 7-2: A snapshot of the streamwise velocity at $t^+ \approx 124$. For demonstration purposes, the lower-bound of the scale range was increased from 0 to 0.5.

log-law region, which begins at $y^+ \approx 30$. As expected, the streamwise and spanwise normal components of the Reynolds stress tensor R^+ feature largest slopes very close to wall, while R^+_{22} the smallest. All the normal components are even functions of y, while the only shear component, R^+_{12}, is an odd function of y. This behavior is a direct consequence of no-slip boundary conditions and can be analytically proven using the Taylor series expansion of velocity components close the wall [140].

7.2 Tangent dynamics of turbulent channel flows

Recall that the simplified computation of sensitivities through Algorithm 3 requires solving a single inhomogeneous tangent equation, which represents the parametric perturbation of the state along a trajectory. To eliminate the effect of the unstable subspace, we recursively "regularize" the tangent solution through orthogonal projection of the unstable component. This process is possible due to the knowledge of the unstable basis, which requires solving m homogeneous tangents at every point on the attractor. In our notation, m is an estimate of the total number of positive Lyapunov exponents (LEs). Figure 7-4 illustrates snapshots of select tangent solutions of the

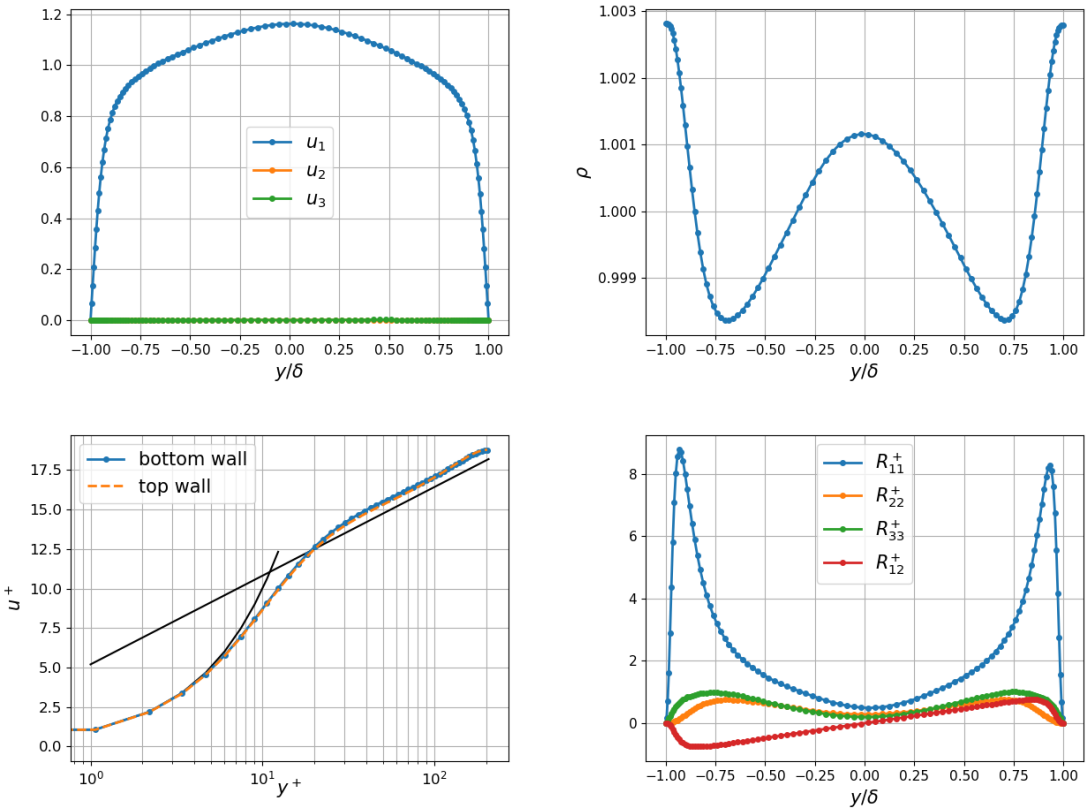

Figure 7-3: Upper plots: wall-normal profiles of the three velocity components (left) and density (right) obtained through averaging in all statistically homogeneous directions (x, z, t). Lower left: u–y relation in "+" units plotted on a semi-logarithmic coordinate system. Solid lines represent the linear function $u^+ = y^+$ and logarithmic function $u^+ = 1/\kappa \log(y^+) + B$, $\kappa = 0.41$, $B = 5.2$. Lower right: non-zero components of the Reynolds stress tensor.

turbulent channel flow. These plots correspond to the test case plotted in Figures 7-2–7-3 and represent the streamwise component only. In this particular case, the dimension of the unstable subspace is $m \approx 260$.

The leading Lyapunov vector (1st LV) features a high level of disturbances close to the walls and almost unperturbed behavior at the center of the channel. These active (non-zero) structures have a streaky elongated shape, which was also observed in a detailed study of the 1st LV in [131]. That paper argued the structure of the leading CV is a direct consequence of the inhomogeneity of low-frequency components of the flow. Our results and the discussion in [131] also agree on the fact that the coherent structures of the leading LV are significantly smaller compared to the primal solution

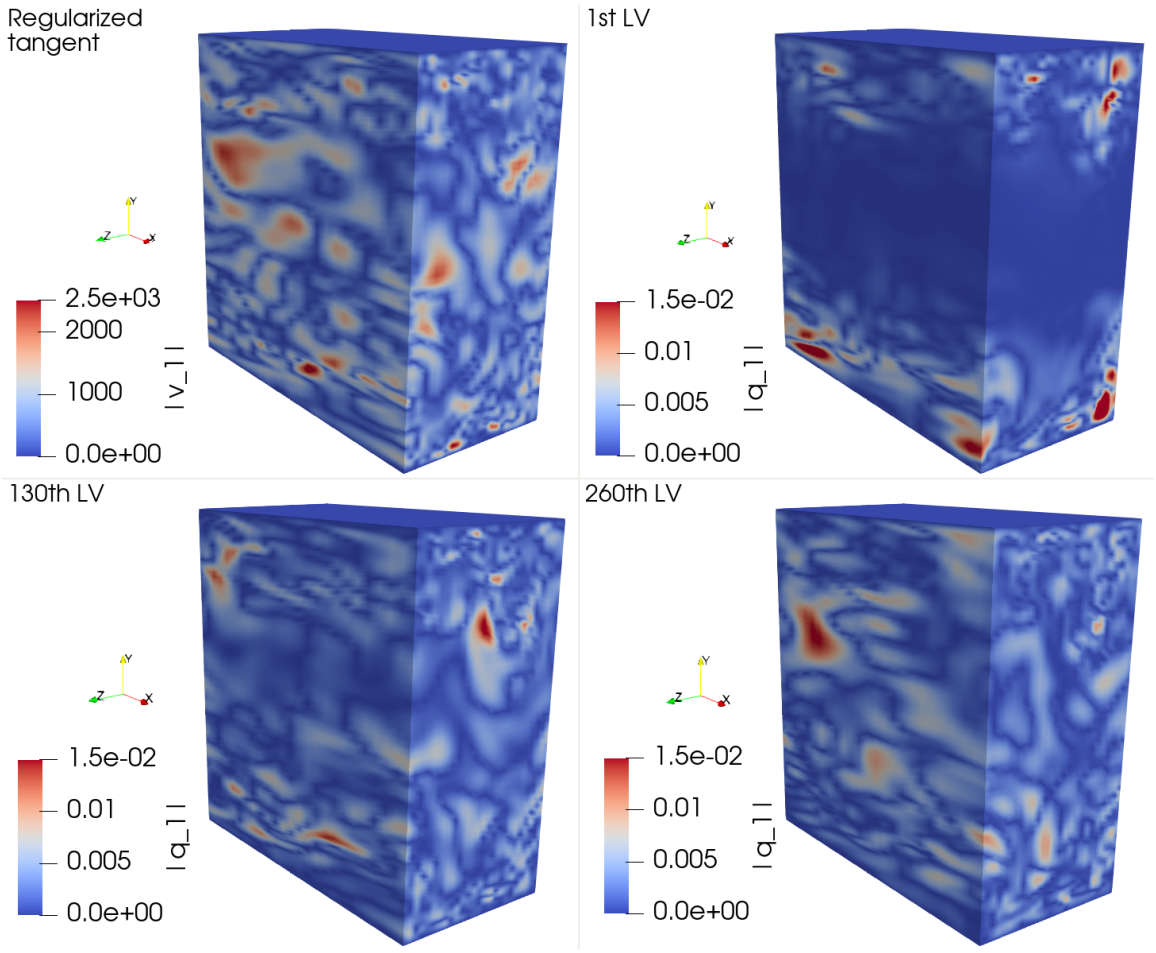

Figure 7-4: Streamwise tangent solutions of the turbulent channel flow at the time instance $t^+ \approx 17.3$. i-th LV represents the i-th most expansive orthonormalized (Gram-Schmidt) Lyapunov vector.

(see Figure 7-2). We also observe that the activity regions tend to move towards the center as the LV index increases. The structure of the 260th CV is no longer localized in space and includes some broader streaks. One can identify some analogies of the geometric features between the LVs plotted in Figure 7-4 and the LVs of the Kuramoto-Sivashinsky equation illustrated in Figure 6-17. These similarities include the wall effect, localization patterns and respective streak sizes. Recall that all the LVs except the leading one lose their physical properties due to the orthonormalization procedure applied in our algorithm, although the spans of the i most unstable Gram-Schmidt vectors and i most unstable covariant Lyapunov vectors (CLVs) are the same. Due to their conserved covariance property, unstable CLVs with smaller positive exponents

might be provide some information on the evolution of secondary instabilities [126]. The structure of the regularized tangent resembles the LVs corresponding to smaller exponents. It typically features large and round high-magnitude streaks close to the center, and several smaller steaks near the wall. The ergodic average of that quantity contracted against an objective function approximates the desired linear response. Therefore, the red regions indicate high-sensitivity regions at a given time instance.

For both full and reduced space-splitting algorithms, it is crucial to have a sufficiently accurate estimate of the number of positive LEs, m. If that parameter is underestimated, the magnitude of the inhomogeneous tangent will rapidly blow-up. Based on the numerical evidence presented in Sections 6.4.2–6.4.3, subtle overestimation of m should not significantly impact the sensitivity approximation. The dependence of m on simulation parameters, e.g., the bulk/friction Reynolds number or characteristic mesh size, is one of the major questions in the turbulence theory. The value of m does not only provide us the critical knowledge on the tangent-dynamic landscape of the flow, but also determines the computational cost of the linear response algorithm.

We shall now look into some aspects of the Lyapunov spectra of the turbulent channel flow described in the previous section. They are computed using Benettin's algorithm [15], which is a part of our reduced space-splitting method (its output are "LEs" in 3). Figure 7-5 illustrates positive parts of the Lyapunov spectra generated for the case described in the previous section. The only parameter that we vary here is the kinematic viscosity. We observe that the leading exponents, usually denoted by λ_1, grow with the Reynolds number. Indeed, the largest LE represents the rate of separation (stretching rates) of the smallest eddies captured by the current mesh. It could also be understood as the reciprocal of the predictability time [94], i.e., time needed to observe sufficiently large deviation of the trajectories of two nearby particles. The localization of the most unstable, i.e., the least predictable, flow regions is represented by the 1st LV shown in Figure 7-4. Less unstable modes corresponding to smaller LEs shown in Figure 7-5 would require generating other CLVs [65], which is beyond the scope of this chapter. All the four curves plotted in this figure feature monotonically decreasing slopes with respect to the index. This result is

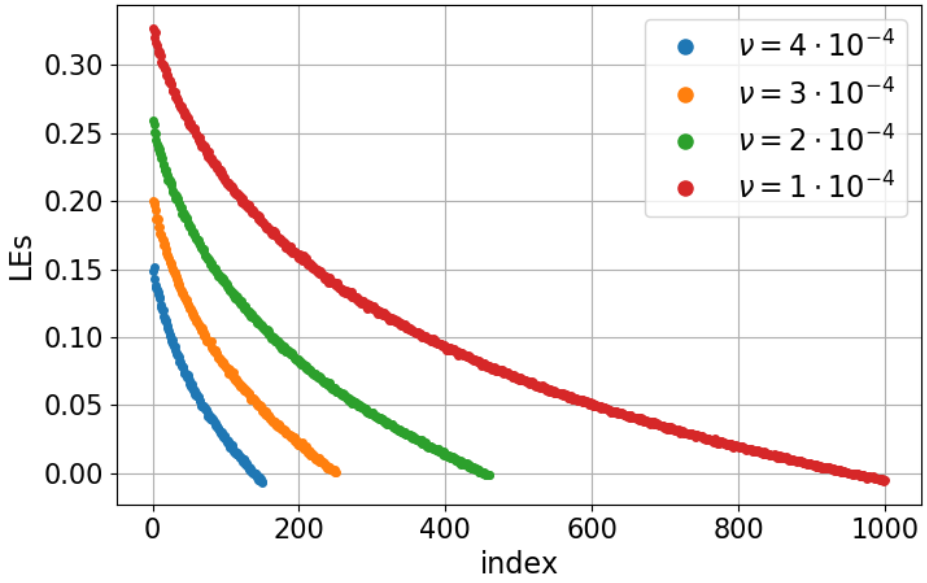

Figure 7-5: Lyapunov spectra of the turbulent channel flow at different values of the kinematic viscosity ν. The remaining input data, including the mesh and initial condition, is fixed. Particular LE values represent ergodic averages over $t^+ \approx 30 - 40$.

consistent with the unstable Lyapunov spectra of spatiotemporal chaos presented in the literature [20, 50]. At any fixed index, the slope of the spectrum curve nonlinearly decreases with ν. Consequently, m grows rapidly with Re_{bulk}, from $m \approx 145$ at $Re_{bulk} = 2,500$ to $m \approx 260$ at $Re_{bulk} = 3,333$ to $m \approx 460$ at $Re_{bulk} = 5,000$ to $m \approx 960$ at $Re_{bulk} = 10,000$.

In the next experiment, we shall study the effect of the mesh resolution on the leading Lyapunov exponent. To maintain critical physical properties of a turbulent flow, e.g., local dissipation rate, logarithmic law of the wall, presence of intense wall-shear stress events and intense enstrophy fluctuations, a sufficiently refined mesh is required [132, 191, 192]. We shall focus on the case with $Re_{bulk} = 10,000$ ($Re_\tau \approx 550$) and use six different meshes whose parameters are summarized in Table 7.1. The total number of degrees of freedom (DOFs) of the underlying dynamical systems is $n = 4\, n_1\, n_2\, n_3$. We plot the time evolution of the ergodic-averages approximating λ_1 using all six meshes in Figure 7-6 Regardless of the mesh used, all the six $\lambda_1(t^+)$ (Lyapunov) curves require about $t^+ \approx 5 - 10$ to reach their plateaus. At $t^+ \approx< 0.1$,

Table 7.1: Mesh parameters of the channel flow at $Re_\tau \approx 550$.

Mesh	(n_1, n_2, n_3)	Δx^+	$\Delta y^+\|_{wall}$	$\Delta y^+\|_{center}$	Δz^+
M1	$(25, 100, 25)$	45.8	2.8	19.6	22.9
M2	$(40, 100, 40)$	28.2	1.7	22.9	14.1
M3	$(55, 110, 55)$	20.4	1.0	23.3	10.2
M4	$(80, 200, 80)$	13.9	0.6	12.7	7.0
M5	$(110, 300, 110)$	10.1	0.6	7.6	5.0
M6	$(110, 400, 110)$	10.1	0.3	6.3	5.0

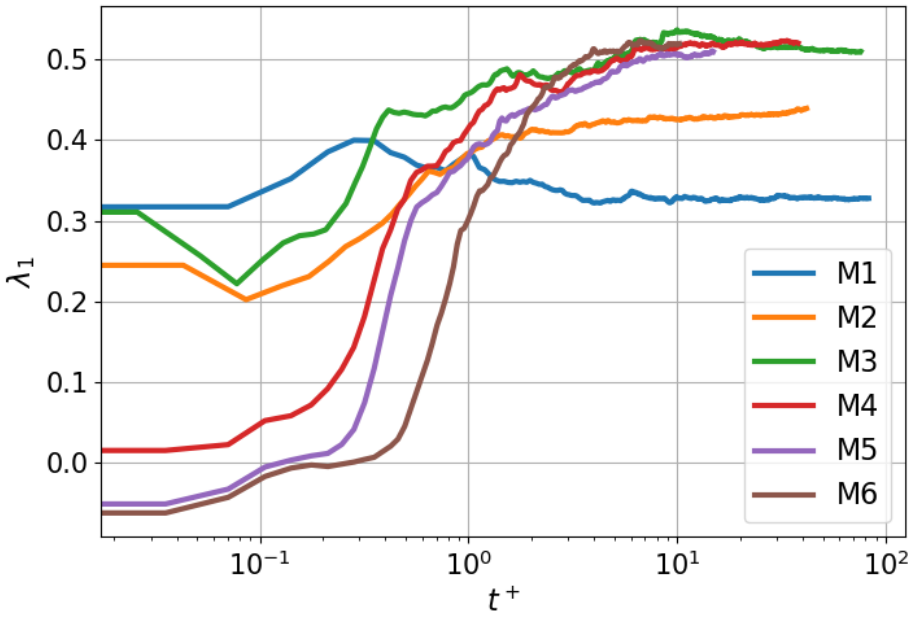

Figure 7-6: Approximations of the leading LE, λ_1, at $Re_\tau \approx 550$.

the M1, M2 and M3 curves are located around the value of 0.3, while the remaining three are closer to zero. This is a consequence of different initial conditions used for those two simulation groups. Recall that the choice of an initial condition does not affect the long-time averages in ergodic systems. As the simulation advances in time, the M1 curve stabilizes at 0.33, M2 around 0.43, while the other four oscillate around the value of 0.51. This experiment indicates that M3 is the optimal mesh for this particular case and no further refinements are necessary to accurately approximate λ_1. Given $m \approx 960$ for M1, by shifting the red curve from Figure 7-5 upwards, we estimate that $m \approx 2,000$ at $Re_\tau \approx 550$ if the mesh is sufficiently resolved. However,

this rough estimate requires further studies as the influence of the mesh size on the slope of Lyapunov curves is still unknown.

7.3 Performance of the reduced linear response algorithm

In the final numerical experiment, we shall investigate the performance of the reduced response algorithm proposed in Chapter 6. The sensitivity approximations will be compared against the finite difference solution obtained by running several forward simulations in the neighborhood of the chosen point on the parametric space. All the results presented here were generated using the same mesh. In particular, $(n_1, n_2, n_3) = (25, 100, 25)$ for all Re_τ considered in this section. Consequently, the grid might be under-resolved in the vicinity of the wall for a larger Re_τ (see Table 7.1).

First, we compute the sensitivity of the long-time average of the volume-averaged kinetic energy K with respect to ν, where

$$K = \frac{1}{L_1\, L_2\, L_3} \int_0^{L_3} \int_{-L_2/2}^{L_2/2} \int_0^{L_1} \rho \left(u_1^2 + u_2^2 + u_3^2\right) dx\, dy\, dz. \tag{7.6}$$

Figure 7-7 illustrates long-time averages of K computed at 270 different values of ν, as well as sensitivities at three select values of ν obtained through the reduced algorithm and cubic regression. Based on several preparatory simulations, the reduced algorithm computes $m = 145$ homogeneous tangents at $\nu = 4 \cdot 10^{-4}$, $m = 460$ at $\nu = 2 \cdot 10^{-4}$ and $m = 960$ at $\nu = 1 \cdot 10^{-4}$. These values of m are our best estimates of the dimension of unstable subspaces for the chosen mesh.

Similarly to the Lyapunov exponents, the sensitivity curves require approximately 10 time plus-units to reach their plateaus. Even at $t^+ > 10$, these curves still feature wavy behavior, especially at higher ν, which is a consequence of long time scales. The upper plot of 7-7 clearly indicates that the statistical error increases with ν. Slow convergence of a similar ergodic averaging process was also observed in the

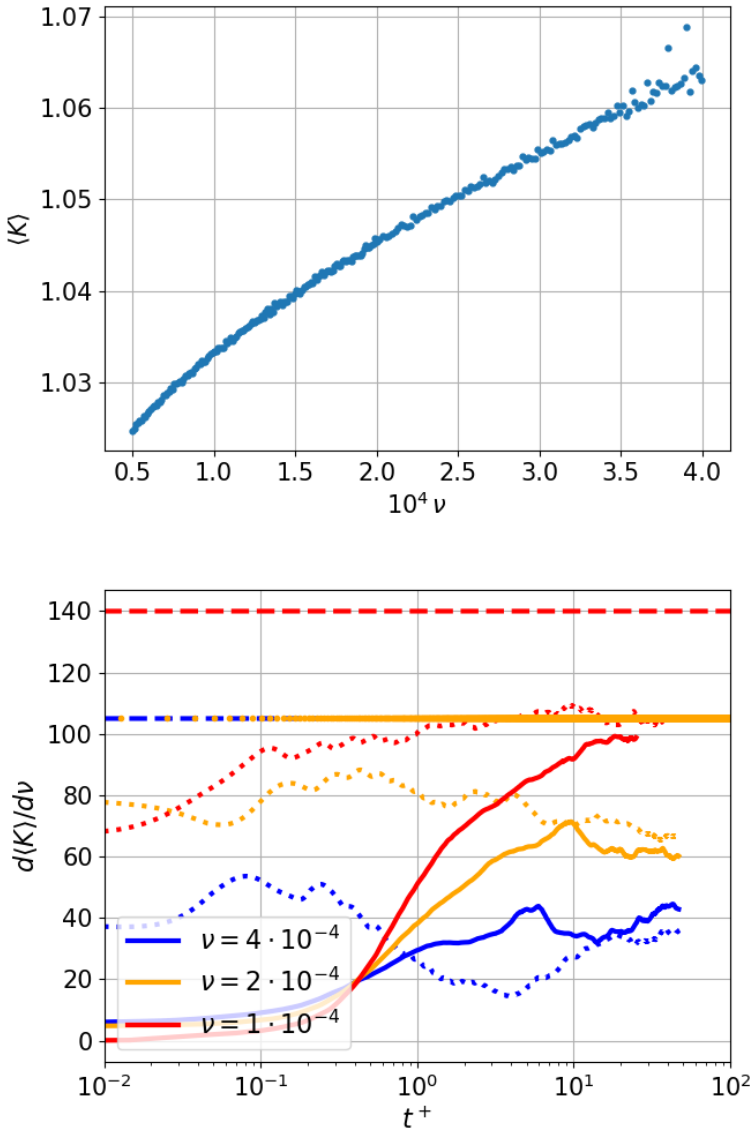

Figure 7-7: Upper plot: long-time averages of K at 270 uniformly distributed values of ν. Every data point represents an independent simulation that ran over $t^+ \approx 124$ time units. Lower plot: solid and dotted curves represent two independent time evolutions of sensitivity approximations obtained though the reduced algorithm at a fixed parameter value. Simulations represented by these two curve types respectively begin statistical averaging after $t^+ \approx 0.05$ and $t^+ \approx 5.5$ of the primal evolution. The dashed horizontal lines indicate reference solutions, which were obtained by taking the derivative of the cubic fit of the data from the upper plot.

recent study [20] in which the shadowing method was applied to a minimal channel flow at $Re_\tau \approx 140$. Here, we observe that the relative error significantly improves as we increase the Reynolds number, from about $60 - 65\%$ at $Re_{bulk} = 2,500$ to $38 - 45\%$ at $Re_{bulk} = 5,000$ to $25 - 30\%$ at $Re_{bulk} = 10,000$. This result is consistent with a major conclusion of Chapter 6. Indeed, the more spatially homogeneous the flow, the better performance of the stable contribution. While the channel flow is spatially homogeneous in both x and z, it features considerable statistical variations in the wall-normal direction. We have already observed a similar behavior around the chaotic-to-laminar transition in solutions of both the Lorenz 63 and Kuramoto-Sivashinsky systems (see Sections 6.4.2–6.4.3). At the same time, we acknowledge that the impact of mesh resolution should be further studied. Indeed, it is still uncertain whether the grid refinement is supportive or obstructive in the context of sensitivity analysis of spatially-extended dynamical systems.

Finally, we shall focus on the case with $Re_{bulk} = 10,000$ and compute the sensitivity of the mean streamwise velocity profile defined as

$$U(y,t) = \frac{1}{L_1\, L_3} \int_0^{L_3} \int_0^{L_1} u(x,y,z,t)\, dx\, dz. \tag{7.7}$$

One of the goals here is to provide additional insights into error sources and its distribution along the wall-normal direction. Notice also that if $J := U$, then the sensitivity approximated as $d\langle J\rangle/d\nu \approx DJ \cdot v$ represents the long-time average of the wall-normal profile of the regularized tangent solution. Results of this test are included in Figure 7-8. The upper plot illustrates the sensitivity approximations obtained using the reduced algorithm and linear regression. The lower one includes statistics-parameter relations at two different points along y: one corresponds to the center of the channel where the relative error is small, while the other one represents the near-wall region characterized by larger errors.

The relative error of the sensitivity approximation, computed with respect to the largest-in-magnitude reference value of $d\langle U\rangle/d\nu$, is quite small. In particular, it does not exceed the value of 5%, except the thin near-wall regions where it reaches the value

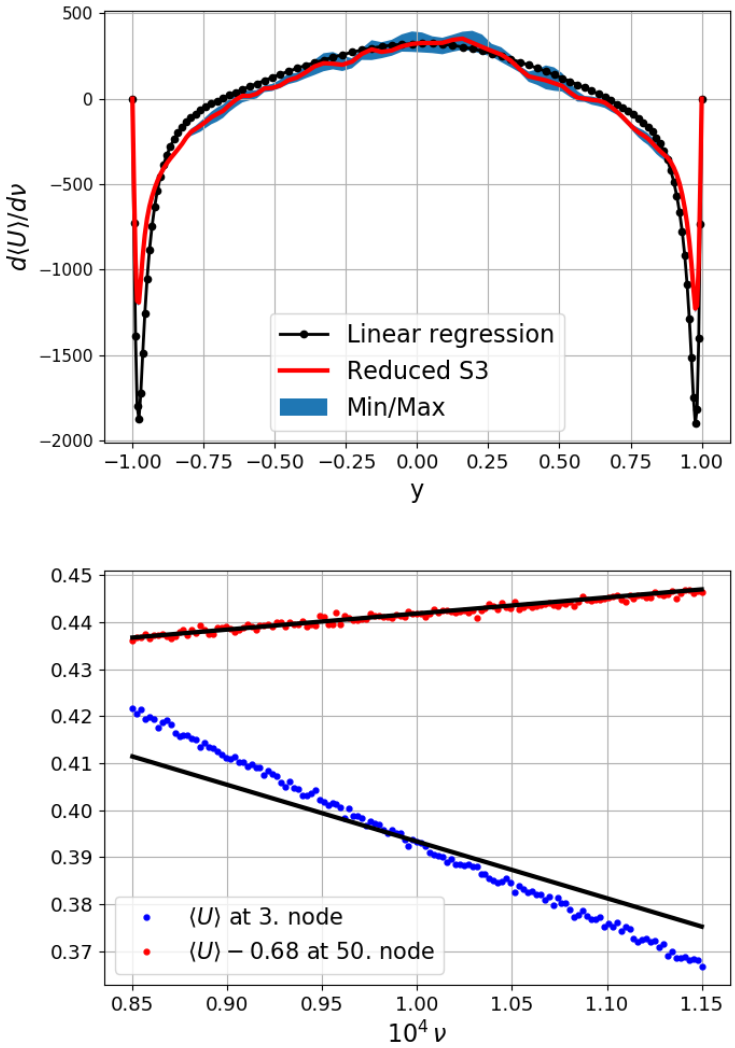

Figure 7-8: Upper plot: sensitivity of the mean streamwise velocity profile at $\nu = 10^{-4}$. The solid red line represents the mean solution of the data obtained between $t^+ \approx 7$ and $t^+ \approx 17$, while the blue shade indicates the region constrained between the minimum and maximum value obtained in that time interval. The reference solution (black) was obtained through linear regression of 135 data points between $\nu = 0.85 \cdot 10^{-4}$ and $\nu = 1.15 \cdot 10^{-4}$ as shown below. Lower plot: long-time averages of the mean streamwise velocity at the 3rd and 50th (center) node from the wall. Every data point represents an independent simulation that ran over $t^+ \approx 124$ time units. The solid black lines represent the corresponding slopes at $\nu = 10^{-4}$ obtained with the reduced S3 method.

of approximately 30%. The blue shade is visible only at the center of the channel, which means that the results obtained at small y^+ are not expected to change visibly. Notice that the reference solution obtained through linear regression can be treated as “sufficiently converged" given the small variance of the time averages computed in the parametric space. The solid black curve corresponding to the 50th (center) node almost ideally fits into the scattered data points. This is no longer true at the 3rd node. We also highlight the fact that the largest discrepancy (error $> 10\%$) between the S3 and linear regression results is located between the 2nd and 10th nodes, which spans the distance $y^+ \in [5, 20]$ from the wall. According to [140], this sub-region is usually referred to as the buffer layer. Notice also that the small error at the 2nd node suggests that our algorithm is expected to be accurate if the wall shear stress is our quantity of interest, i.e., when $J := \tau_w = \rho\,\nu\,dU/dy$.

7.4 Summary

Turbulent flows are prevalent in nature and engineering systems. Understanding the dynamics of such flows is critical in many aspects of our life such as the weather prediction, control of hydromechanical systems, design of advanced aerospace technology. In this chapter, we considered one of the simplest turbulent flows, i.e., nearly incompressible isothermal channel flows. Advancement of reliable numerical tools for the applied analysis of channel units is highly desired by the scientific and engineering communities. These tools could be utilized to better understand the effect of boundary layer, nature of dynamical instabilities, extreme events, potential singularity formation and several other fundamental topics. From the mathematical point of view, one of the inherent properties of turbulence is chaoticity that we defined in several ways throughout this thesis. Thus, the linear response machinery that we developed in the previous chapters could potentially be used in the comprehensive sensitivity analysis.

In this chapter, we studied some aspects of the tangent dynamics of channel flows. In our test case, we observed that the number of positive Lyapunov exponents grows

10-fold when the Reynolds number increases from $Re_{bulk} \approx 2,500$ to $Re_{bulk} \approx 10,000$. This nonlinear relation sheds some light on the total cost of the linear response algorithm when applied to highly-turbulent cases. We conclude that even geometrically simple cases might reach $m = 1,000$ or higher at moderate Reynolds numbers, which means that high performance computing resources are a must in this type of study. Recall that the cost of the reduced S3 algorithm is proportional to m^2. We also observed that the leading Lyapunov exponent can be dramatically underestimated if the mesh is under-resolved. This deficiency does not affect the algorithm *per se*, but it might produce results violating fundamental physics of the flow. The leading LE determines the maximum rate of separation of two nearby trajectories of a dynamical system. By analyzing the structure of the corresponding Lyapunov vector, one can identify instability regions across the spatial domain. We observed that the most expansive regions are localized and concentrated close to the wall, which is consistent with the literature. The structure of the regularized tangent solution is completely different. It contains both smaller, elongated streaks high-magnitude at the wall and large broad high-magnitude structures close to the center of the channel.

We applied the reduced S3 algorithm to compute sensitivities of both the kinematic energy and streamwise velocity profile with respect to the kinematic viscosity. Our results were compared against reference values obtained through polynomial regression. The latter were obtained by running hundreds of forward simulations at different parameter values. In all test cases, our method consistently provided underestimated approximations of sought-after sensitivities. However, we noticed significant decrease of the relative error from about 60% at $Re_{bulk} \approx 2,500$ to $25 - 30\%$ at $Re_{bulk} \approx 10,000$. This result is consistent with our results presented in Chapter 6. Indeed, as we approach the laminar-to-turbulence transition, the effect of the ergodic measure gradient contributes more to the total sensitivity. In our study of the velocity profile, we observed good agreement between the S3 and regression approximations. The relative errors, computed with respect to the maximum reference value across the wall-normal direction, did not exceed a few percent with the exception of the buffer layers, where the errors reached the value of about 30% at $Re_{bulk} \approx 10,000$.

In general, the performance of the reduced S3 method is similar to its main competitors, i.e., the least-squares shadowing (LSS) approach and its derivatives. Our algorithm is significantly simpler compared to shadowing methods, as it involves solving a single inhomogeneous tangent that is recursively regularized by orthogonally projecting out its unstable component. However, the leading term of the flop count remains the same, because all of these methods require the knowledge on the unstable directions along the random trajectory. To the best of our knowledge, most of popular methods rely on the QR factorization as a part of the process.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 8

Final conclusions and future work

This thesis addresses several topics at the intersection of dynamical systems, linear response theory, numerical analysis and chaotic physics. Its main product is the generalized space-splitting (S3) method for sensitivity analysis of hyperbolic dynamical systems with $m \geq 1$ positive Lyapunov exponents. S3 can be viewed as a one-step Monte Carlo procedure that collects and averages data along random trajectories. The crux of the S3 method is to regularize Ruelle's original expression for the linear response, because its direct evaluation is infeasible due to the phenomenon of exploding Jacobian products, i.e., the *butterfly effect*. The reformulation of Ruelle's result involves two major steps: linear splitting of the input perturbation and integration by parts. The purpose of splitting is to decompose the input such that all the resulting components are computable through tangent equations that are immune to the butterfly effect. This can be achieved by representing one of the terms as a linear combination of an orthonormal basis of the unstable subspace. That term is usually referred to as the unstable contribution and effectively involves a sum of directional derivatives that can be partially integrated on unstable manifolds. Moreover, by adjusting the scalar splitting coefficients, the remaining contribution can be solved by evolving an inhomogeneous tangent that is bounded in time.

Conceptually, the most sophisticated part of the proposed algorithm is the procedure for differentiating Sinai-Ruelle-Bowen (SRB) measures. It assumes the smoothness of the limiting ergodic distributions and requires the knowledge of the unstable

manifold curvature represented by second-order derivatives of the coordinate chart. Consequently, one needs to evolve $\mathcal{O}(m^2)$ second-order tangent equations in time to differentiate the SRB measure. Quantile functions are convenient in parameterizing smooth manifolds, as they facilitate a closed-form and linearly invariant representation of the logarithmic density gradient in terms of derivatives of the coordinate charts. The property of invariance to linear transformations allows for consistent stabilization of the tangent recursions through step-by-step re-orthonormalization of the chart gradient. Using this manifold description, one can recursively differentiate the scalar coefficients arising in the splitting process, which is necessary to complete the S3 algorithm. Under the assumption of *uniform hyperbolicity*, all these recursions rigorously converge along a random trajectory at an exponential rate. We conclude that the total cost of the space-split algorithm applied to n-dimensional systems is $\mathcal{O}(n^3\, m^2)$ per time step. In practice, dynamical systems feature a sparse structure as they are derived by converting PDEs to ODEs using local-stencil methods, such as the finite element method. In addition, the number of positive LEs is typically orders of magnitude lower than the system's dimension, in which case the total cost of S3 can be reduced to $\mathcal{O}(n\, m^3)$ per time step. Extending the S3 method to continuous-time systems requires further, yet minor, modifications. In particular, the original unstable-stable splitting should also include an extra term that is aligned with the flow. While this upgrade guarantees constant variance of the regularized tangent, a Schur complement system must be solved because the evolution vector is generally non-orthogonal to the unstable basis. With this adjustment, the leading term of the flop count does not change.

Several physical dynamical systems, e.g., the Kuramoto-Sivashinsky or Navier-Stokes equation, are non-hyperbolic, which implies that there are local tangencies between tangent subspaces, and the smoothness of the SRB measure is no longer guaranteed. It also implies that the computation of unstable contributions by differentiating SRB measures in m most expansive orthogonal directions might be infeasible. However, in certain systems, the unstable part of the sensitivity could be neglected. Through double partial integration of Ruelle's original expression, one

can show that the effect of measure gradients is negligible if the objective function is aligned with the most expansive direction of the tangent space. This observation might have important implications for large physical systems featuring statistical homogeneity in space. Indeed, it is possible to represent any objective function in an infinite number of ways if the long-time statistics of individual spatial coordinates is approximately constant. The dimension of the unstable subspace is typically orders of magnitude smaller than the system's dimension. Consequently, by increasing the number of degrees of freedom, the likelihood of the desired alignment grows as well. Our numerical experiments applied to Lorenz 96, Kuramoto-Sivashinsky and turbulent channel flow clearly confirm our conjecture. Therefore, the linear response could be accurately approximated by computing a single inhomogeneous tangent that is orthogonal to the unstable subspace. Eliminating the unstable part from the S3 algorithm reduces the total cost to $\mathcal{O}(n\,m^2)$ per time step. The m^2 factor is associated with the QR factorization that is required for the orthogonalization of unstable bases along random orbits. We demonstrate that m could be $\mathcal{O}(10^3)$ in fairly simple turbulent channel flows with moderate Reynolds numbers. Thus, the use of high performance computing hardware is a must in comprehensive sensitivity analysis of turbulent flows.

A number of unanswered questions still remain, which could serve as a starting point of future work. The following list outlines three potentially impactful research directions and, to the best of our knowledge, have not been explored yet:

- **Development of the correction term for sensitivity analysis of transitional chaos.** Understanding the laminar-turbulent transition is critical in the design of hypersonic aircraft, blood-contacting medical devices, and several other applications. Computing the linear response of the underlying dynamical systems is difficult, because the dimensionality of tangent subspaces is ultrasensitive to input parameters. In this case, both the effect of the ergodic measure change and the curvature of unstable manifolds are significant. That fact is neglected by shadowing-based techniques (including our reduced S3 method) and, consequently, they are all doomed to fail when applied to transitional flows.

Since the shadowing contribution is only a part of the full response of a chaotic system, the goal is to improve the existing methods by deriving and implementing a "correction" term that will account for the neglected part. This theme is full of fundamental and algorithmic challenges. Local non-hyperbolic behavior, which is quite common in real-world chaotic systems, might be a source of rough measure distributions leading to numerical instabilities. On the practical side, one requires Hessian operators to compute the curvature of manifolds. Their efficient computation is the major, yet still unresolved, step in bridging the linear response theory to engineering.

- **Global sensitivity analysis.** In this thesis, we propose numerical procedures for sensitivity analysis at a given point in the parametric space. Indeed, Ruelle derived an expression enabling local analysis. It is not uncommon that the response or long-time statistics of a physical system is strongly oscillatory (non-convex). This could happen, for example, when the statistics-parameter relation is inherently rough (see the example of Lorenz 63 in Chapter 4) and/or violent physical phenomena affecting spatial homogeneity occur (see [6]). In such cases, we are more interested in the expected value of sensitivity over a certain parametric subdomain rather than pointwise parametric derivatives. Several different approaches addressing this issue could be explored in detail, e.g., adding smoothening noise to the governing equations and/or blending tangent solutions corresponding to different parameter instances.

- **Integration of the S3 method with LES models and generic ROM techniques.**

  The purpose of this research theme is to combine the developed linear response machinery with methods reducing the dimensionality of complex physical systems. We have observed that the performance of the reduced S3 method applied to turbulent channel flows dramatically improves as the Reynolds number increases. Thus, it could be worth investigating its performance at Reynolds numbers an order of magnitude larger than those considered in this work. Given

the rapid growth of the number of positive LEs, a natural choice would be to incorporate certain physics-based models, e.g., large eddy simulation (LES), dramatically alleviating the total cost. Another fundamental question concerns the influence of model order reduction techniques on long-time statistics and tangent landscape of the underlying dynamical systems. Successful integration of S3 and reduced order modeling (ROM) techniques could serve as a powerful and affordable tool for applied analysis of large-scale systems.

THIS PAGE INTENTIONALLY LEFT BLANK

Bibliography

[1] R. V. Abramov and A. J. Majda. Blended response algorithms for linear fluctuation-dissipation for complex nonlinear dynamical systems. *Nonlinearity*, 20(2793), 2007.

[2] R. V. Abramov and A. J. Majda. New Approximations and Tests of Linear Fluctuation-Response for Chaotic Nonlinear Forced-Dissipative Dynamical Systems. *Journal of Nonlinear Science*, 18:303–341, 2008.

[3] M. Aengenheyster, Q. Y. Feng, F. Van der Ploeg, and H. A. Dijkstra. The point of no return for climate action: effects of climate uncertainty and risk tolerance. *Earth System Dynamics*, 9:1085–1095, 2018.

[4] J. Alves, C. Bonatti, and M. Viana. SRB measures for partially hyperbolic systems whose central direction is mostly expanding. *Inventiones Mathematicae*, 140:351–398, 2000.

[5] F. T. Arecchi, G. Giacomelli, A. Lapucci, and R. Meucci. Two-dimensional representation of a delayed dynamical system. *Physical Review A*, 45, 1992.

[6] W. Arias-Ramírez. *A multi-fidelity approach to sensitivity estimation in large eddy simulations*. PhD thesis, University of Maryland, 2022.

[7] L. Arnold. *The Multiplicative Ergodic Theorem on Bundles and Manifolds. In: Random Dynamical Systems*. Springer, Berlin, Heidelberg, 1998.

[8] V. I. Arnold and A. Avez. *Ergodic problems of classical mechanics*, pages 1–286. W. A. Benjamin, 1968.

[9] L. Arriola and J. M. Hyman. Sensitivity analysis for uncertainty quantification in mathematical models. In G. Chowell, J. M. Hyman, L. M. A. Bettencourt, and C. Castillo-Chavez, editors, *Mathematical and Statistical Estimation Approaches in Epidemiology*, pages 195–247. Springer Netherlands, 2009.

[10] V. Baladi. *Positive Transfer Operators and Decay of Correlations*. World Scientific, 2000.

[11] V. Baladi. Linear response, or else. *arXiv e-prints*, arXiv:1408.2937, 2014.

[12] V. Baladi and D. Smania. Linear response formula for piecewise expanding unimodal maps. *Nonlinearity*, 21:677–711, 2008.

[13] D. Balagangadhar and R. Subrata. Design sensitivity analysis and optimization of steady fluid-thermal systems. *Computer Methods in Applied Mechanics and Engineering*, 190:5465–5479, 2001.

[14] A. Ben-Israel and R. Gilbert. *Fundamental theorem of calculus*, pages 322–343. Springer Vienna, Vienna, 2002.

[15] G. Benettin, L. Galgani, A. Giorgilli, and J.-M. Strelcyn. Lyapunov Characteristic Exponents for smooth dynamical systems and for hamiltonian systems; A method for computing all of them. Part 2: Numerical application. *Meccanica*, 15:21–30, 1980.

[16] K. K. Benke, K.E. Lowell, and A.J. Hamilton. Parameter uncertainty, sensitivity analysis and prediction error in a water-balance hydrological model. *Mathematical and Computer Modelling*, 47(11-12):1134–1149, 2008.

[17] P. B. Bertsekas and J. N Tsitsiklis. *Derived Distributions*, pages 202–216. Athena Scientific, Nashua, NH, USA, 2008.

[18] P. Blonigan. *Least Squares Shadowing for Sensitivity Analysis of Large Chaotic Systems and Fluid Flows*. PhD thesis, Massachusetts Institute of Technology, 2016.

[19] P. Blonigan and Q. Wang. Probability density adjoint for sensitivity analysis of the Mean of Chaos. *Journal of Computational Physics*, 270:660–686, 2014.

[20] P. J. Blonigan. Adjoint sensitivity analysis of chaotic dynamical systems with non-intrusive least squares shadowing. *Journal of Computational Physics*, 348:803–826, 2017.

[21] P. J. Blonigan and Q. Wang. Least squares shadowing sensitivity analysis of a modified Kuramoto–Sivashinsky equation. *Chaos, Solitons & Fractals*, 64:16–25, 2014.

[22] T. Bodai, V. Lucarini, and F. Lunkeit. Can we use linear response theory to assess geoengineering strategies? *arXiv e-prints*, arXiv:1803.09606, 2020.

[23] G. Boffetta, G. Lacoratta, S. Musacchio, and A. Vulpiani. Relaxation of finite perturbations: Beyond the fluctuation-response relation. *Chaos*, 13, 2003.

[24] C. Bonatti, L. J. Díaz, and M. Viana. *Dynamics beyond uniform hyperbolicity*. Springer, Berlin, Heidelberg, 2005.

[25] K. Burns, D. Dolgopyat, Y. Pesin, and M. Pollicott. Stable ergodicity for partially hyperbolic attractors with negative central exponents. *Journal of Modern Dynamics*, 2:63–81, 2008.

[26] A. Caicedo-Casso, H.-W. Kang, S. Lim, and C. I. Hong. Robustness and period sensitivity analysis of minimal models for biochemical oscillators. *Science Reports*, 5(13161), 2015.

[27] Y. Cao, S. Li, L. Petzold, and R. Serban. Adjoint sensitivity analysis for differential-algebraic equations: The adjoint DAE system and its numerical solution. *SIAM Journal of Scientific Computing*, 24(3):1076–1089, 2003.

[28] N. Chandramoorthy. *An efficient algorithm for sensitivity analysis of chaotic systems*. PhD thesis, Massachusetts Institute of Technology, 2021.

[29] N. Chandramoorthy, P. Fernandez, C. Talnikar, and Q. Wang. Feasibility analysis of ensemble sensitivity computation in turbulent flows. *AIAA Journal*, 57(10), 2019.

[30] N. Chandramoorthy, L. Magri, and Q. Wang. Variational optimization and data assimilation in chaotic time-delayed systems with automatic-differentiated shadowing sensitivity. *arXiv e-prints*, arXiv:2011.08794, 2020.

[31] N. Chandramoorthy and Q.Wang. Efficient computation of linear response of chaotic attractors with one-dimensional unstable manifolds. *SIAM Journal on Applied Dynamical Systems*, 21(2):735–781, 2022.

[32] N. Chandramoorthy and Q. Wang. A computable realization of Ruelle's formula for linear response of statistics in chaotic systems. *arXiv e-prints*, arXiv:2002.04117, 2020.

[33] N. Chandramoorthy and Q. Wang. An ergodic-averaging method to differentiate covariant Lyapunov vectors. *Nonlinear Dynamics*, 104:4083–4102, 2021.

[34] N. Chandramoorthy and Q. Wang. On the probability of finding a nonphysical solution through shadowing. *Journal of Computational Physics*, 440(110389), 2021.

[35] J. D. Chekroun, M. C. Neelin, D. Kondrashov, J. C. McWilliams, and M. Ghil. Rough parameter dependence in climate models and the role of Ruelle-Pollicott resonances. *Proceedings of the National Academy of Sciences*, 111:1684–1690, 2014.

[36] N. I. Chernov. Limit theorems and Markov approximations for chaotic dynamical systems. *Probability Theory and Related Fields*, 101:321–362, 1995.

[37] D. Chrysos. An extension of the Filon method for the accurate numerical integration of rapidly varying functions. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 28(11):L373–L377, 1995.

[38] G. A. Chua and Y. Liu. Sensitivity analysis on responsive pricing and production under imperfect demand updating. *Naval Research Logistics*, 66(7):529–546, 2019.

[39] I. Cionni, G. Visconti, and F. Sassi. Fluctuation dissipation theorem in a general circulation model. *Geophysical Research Letters*, 31, 2004.

[40] V. Climenhaga, D. Dolgopyat, and Y. Pesin. Non-stationary non-uniform hyperbolicity: SRB measures for dissipative maps. *Communications in Mathematical Physics*, 346:553–602, 2016.

[41] H. Crimmins and G. Froyland. Fourier approximation of the statistical properties of Anosov maps on tori. *Nonlinearity*, 33(11), 2020.

[42] A. Cruz and P. Varandas. SRB measures for partially hyperbolic attractors of local diffeomorphisms. *Ergodic Theory and Dynamical Systems*, 40:1545–1593, 2020.

[43] D. Cruz-Uribe and C. J. Neugebauer. Sharp error bounds for the trapezoidal rule and Simpson's rule. *Journal of Inequalities in Pure and Applied Mathematics*, 3(49), 2002.

[44] G. De Nittis and M. Lein. *Linear Response Theory: An Analytic-Algebraic Approach*. Springer, Cham, 2017.

[45] J. Ding and T. Y. Li. Markov finite approximation of Frobenius-Perron operator. *Nonlinear Analysis: Theory, Methods & Applications*, 17:759–772, 1991.

[46] D. Dolgopyat. On differentiability of SRB states for partially hyperbolic systems. *Inventiones mathematicae*, 155:389–449, 2004.

[47] D. J. Driebe. *The Bernoulli Map*, pages 29–43. Springer, 1996.

[48] H. A. Dwyer and T. Peterson. Study of turbulent flow with sensitivity analysis. *AIAA Journal*, 19(10):1309–1314, 2020.

[49] D. A. Egolf. Equilibrium regained: From nonequilibrium chaos to statistical mechanics. *Science*, 287:101–104, 2000.

[50] D. A. Egolf, I. V. Melnikov, W. Pesch, and R. E. Ecke. Mechanisms of extensive spatiotemporal chaos in Rayleigh-Bénard convection. *Nature*, 404:733–736, 2000.

[51] V. E. Ershov and A. B. Potapov. On the conceptof stationary Lyapunov basis. *Physica D*, 118:167–198, 1998.

[52] G. Eyink, T. Haine, and D. Lea. Ruelle's linear response formula, ensemble adjoint schemes and lévy flights. *Nonlinearity*, 17:1867, 2004.

[53] J. Faber and H. Koppelaar. Chaos theory and social science: A methodological analysis. *Quality and Quantity*, 28:421–433, 1994.

[54] A. Fasoli, F. Skiff, R. Kleiber, M. Q. Tran, and P. J. Paris. Dynamical chaos of plasma ions in electrostatic waves. *Physical Review Letters*, 70:303–306, 1993.

[55] A. Fichtner, H.-P. Bunge, and H. Igel. The adjoint method in seismology: II. Applications: traveltimes and sensitivity functionals. *Physics of the Earth and Planetary Interiors*, 157:105–123, 2006.

[56] K. J. Fidkowski and D. L. Darmofal. Review of output-based error estimation and mesh adaptation in computational fluid dynamics. *AIAA Journal*, 49:673–694, 2011.

[57] S. Fiori. Neural systems with numerically-matched input–output statistic: variate generation. *Natural Processing Letters*, 23:143–170, 2006.

[58] Y. W. Foo, C. Goh, and Y. Li. Machine learning with sensitivity analysis to determine key factors contributing to energy consumption in cloud data centers. *2016 International Conference on Cloud Computing Research and Innovations (ICCCRI)*, pages 107–113, 2016.

[59] G Galavotti and E. G. D. Cohen. Dynamical ensembles in stationary states. *Journal of Statistical Physics*, 80:931–970, 1995.

[60] P. Gaspard. Diffusion, effusion, and chaotic scattering: An exactly solvable liouvillian dynamics. *Journal of Statistical Physics*, 68:673–747, 1992.

[61] F. Geng, I. Kalkman, A. S. J. Suiker, and B. Blocken. Sensitivity analysis of airfoil aerodynamics during pitching motion at a Reynolds number of $1.35 \cdot 10^5$. *Journal of Wind Engineering and Industrial Aerodynamics*, 183:315–332, 2018.

[62] G. Giacomelli, R. Meucci, A. Politi, and F. T. Arecchi. Defects and spacelike properties of delayed dynamical systems. *Physical Review Letters*, 73, 1994.

[63] M.B. Giles and N.A. Pierce. An introduction to the adjoint approach to design. *Flow, Turbulence and Combustion*, 65:393–415, 2000.

[64] S. N. Giles and S. Fiori. Glomerular filtration rate estimation by a novel numerical binning-less isotonic statistical bivariate numerical modeling method. *Information*, 10(100), 2019.

[65] F. Ginelli, H. Chaté, R. Livi, and A. Politi. Covariant Lyapunov vectors. *Journal of Physics A: Mathematical and Theoretical*, 46(254005), 2013.

[66] J.-M. Ginoux and C. Letellier. Van der Pol and the history of relaxation oscillations: Toward the emergence of a concept. *Chaos*, 22(023120), 2012.

[67] H. B. Gregersen and L. Sailer. Chaos theory and its implications for social science research. *Human Relations*, 46(7):777–802, 1993.

[68] A. Gritsun and G. Branstator. Climate response using a three-dimensional operator based on the fluctuation–dissipation theorem. *Journal of Atmospheric Sciences*, 64:2558–2575, 2016.

[69] A. Gritsun, G. Branstator, and A. Majda. Climate response of linear and quadratic functionals using the Fluctuation–Dissipation Theorem. *Journal of Atmospheric Sciences*, 65:2824—2841, 2008.

[70] A. Gritsun and V. Lucarini. Fluctuations, response, and resonances in a simple atmospheric model. *Physica D: Nonlinear Phenomena*, 349:62–76, 2017.

[71] A. Gronkis, D. Heitz, and E. Mémin. Inflow and initial conditions for di- rect numerical simulation based on adjoint data assimilation. *Journal of Computational Physics*, 242:480–497, 2013.

[72] M. S. Gutiérrez and V. Lucarini. Response and sensitivity using Markov chains. *Journal of Statistical Physics*, 179:1572–1593, 2020.

[73] M. Hairer and A. J. Majda. A simple framework to justify linear response theory. *Nonlinearity*, 23(4), 2010.

[74] J. Hart and B. Van Bloemen Waanders. Enabling hyper-differential sensitivity analysis for ill-posed inverse problems. *arXiv e-prints*, arXiv:2106.11813, 2022.

[75] S. R. Haskey, M. J. Lanctot, Y. Q. Liu, C. Paz-Soldan, J. D. King, B. D. Blackwell, and O. Schmitz. Effects of resistivity and rotation on the linear plasma response to non-axisymmetric magnetic perturbations on diii-d. *Plasma Physics and Controlled Fusion*, 57(2):025015, 2015.

[76] A. Hay, J. Borggaard, and D. Pelletier. Local improvements to reduced-order models using sensitivity analysis of the proper orthogonal decomposition. *Journal of Fluid Mechanics*, 629:41–72, 2009.

[77] E. Henin and I. Prigogine. Entropy, dynamics, and molecular Chaos. *Proceedings of the National Academy of Sciences of the United States of America*, 71(7):2618–2622, 1974.

[78] R. M. Hicks and P. A. Henne. Wing design by numerical optimization. *AIAA Journal*, 15(7), 1978.

[79] C. Hu, X. Yang, X. Zhu, and Z. Du. Stability and structural sensitivity analysis of the turbulent flow in the narrow vaneless diffuser with mean flow method. *Computers & Fluids*, 177:46–57, 2018.

[80] J.-T. Hwang and H. Rabitz. The Green's function method of sensitivity analysis in quantum dynamics. *Journal of Chemical Physics*, 79(4609), 1979.

[81] J. Infante Barbosa, C.M. Mota Soares, and C.A. Mota Soares. Sensitivity analysis and shape optimal design of axisymmetric shell structures. *Computing Systems in Engineering*, 2(5-6):525–533, 1991.

[82] A. Iserles and S. P. Norsett. On quadrature methods for highly oscillatory integrals and their implementation. *BIT Numerical Mathematics*, 44:755–772, 2004.

[83] A. Jameson. Aerodynamic design via control theory. *Journal of Scientific Computing*, 3(3):233–260, 1988.

[84] K. D. Kantarakias and G. Papadakis. Sensitivity analysis of chaotic systems using a frequency-domain shadowing approach. *Journal of Computational Physics*, 474(111757), 2022.

[85] A. Karimi and M. R. Paul. Extensive chaos in the Lorenz-96 model. *Chaos*, 20:043105, 2010.

[86] U. Kirsch. Efficient sensitivity analysis for structural optimization. *Computer Methods in Applied Mechanics and Engineering*, 117:143–156, 1994.

[87] H. Kontani and Y. Yamakawa. Linear response theory for shear modulus C_{66} and raman quadrupole susceptibility: Evidence for nematic orbital fluctuations in fe-based superconductors. *Physical Review Letters*, 113:047001, 2014.

[88] R. Kubo. The fluctuation-dissipation theorem. *Reports on Progress in Physics*, 29(255), 1966.

[89] P. V. Kuptsov. Fast numerical test of hyperbolic chaos. *Physical Review E*, 85:015203, 2012.

[90] P.V. Kuptsov and U. Parlitz. Theory and Computation of Covariant Lyapunov Vectors. *Nonlinear Science*, 22:727–762, 2012.

[91] Y. Kuramoto and T. Tsuzuki. Persistent propagation of concentration waves in dissipative media far from thermal equilibrium. *Progress of Theoretical Physics*, 55:356–369, 1976.

[92] P. K. Kuznetsov. *Hyperbolic chaos: a physicist's view.* Springer, Berlin, Heidelberg, 2012.

[93] S.P. Kuznetsov and E.P. Seleznev. A strange attractor of the smale-williams type in the chaotic dynamics of a physical system. *Journal of Experimental and Theoretical Physics*, 102:355–364, 2006.

[94] J. Labahn, G. Nastac, L. Magri, and M Ihme. Determining dynamic content of turbulent flow LES using the Lyapunov exponent. *Center for Turbulence Research Annual Research Briefs*, pages 77–89, 2017.

[95] P. L. Langen and V. A. Alexeev. Estimating $2xCO_2$ warming in an aquaplanet GCM using the fluctuation-dissipation theorem. *Geophysical Research Letters*, 32, 2016.

[96] J. Larsson. *Grid-adaptation for chaotic multi-scale simulations as a verification-driven inverse problem*, pages 1–17. AIAA, 2018.

[97] J. Larsson and Q. Wang. The prospect of using large eddy and detached eddy simulations in engineering design, and the research required to get there. *Philisophical transactions of Royal Society A*, 372:20130329, 2014.

[98] D. Lasagna. Sensitivity Analysis of Chaotic Systems Using Unstable Periodic Orbits. *SIAM Journal on Applied Dynamical Systems*, 17(1):547–580, 2018.

[99] D. Lasagna, A. Sharma, and J. Meyers. Periodic shadowing sensitivity analysis of chaotic systems. *Journal of Computational Physics*, 391:119–141, 2019.

[100] A. Lasota and J. Yorke. On the existence of invariant measures for piecewise monotonic transformations. *Transactions of the American Mathematical Society*, 186, 1973.

[101] V. Latora, A. Rapisarda, and S. Ruffo. Chaos and statistical mechanics in the Hamiltonian mean field model. *Physica D: Nonlinear Phenomena*, 131:38–54, 1999.

[102] D. J. Lea, M. R. Allen, and T. W. Haine. Sensitivity analysis of the climate of a chaotic system. *Tellus*, 52(5):523–532, 2000.

[103] J. L. Lebowitz and H. Spohn. A Gallavotti–Cohen-type symmetry in the large deviation functional for stochastic dynamics. *Journal of Statistical Physics*, 95:333–365, 1999.

[104] C. E. Leith. Climate response and fluctuation dissipation. *Journal of the Atmospheric Sciences*, 32, 1975.

[105] V. Lembo, V. Lucarini, and F. Ragone. Beyond forcing scenarios: predicting climate change through response operators in a coupled general circulation model. *Scientific Reports*, 10(8668), 2020.

[106] D. Levin. Procedures for computing one- and two-dimensional integrals of functions with rapid irregular oscillations. *Mathematics of Computation*, 38(158):531–538, 1982.

[107] T.-Y. Li. Finite approximation for the Frobenius-Perron operator. A solution to Ulam's conjecture. *Journal of Approximation Theory*, 17, 1976.

[108] H. Lin and C. Wang. Influences of electromagnetic radiation distribution on chaotic dynamics of a neural network. *Applied Mathematics and Computation*, 369(124840), 2020.

[109] C. Liverani. Decay of correlations for piecewise expanding maps. *Journal of Statistical Physics*, 78:1111–1129, 1995.

[110] E. Lorenz. Deterministic nonperiodic flow. *Journal of Atmospheric Sciences*, 32(10):2022–2026, 1963.

[111] E. Lorenz. *Predictability – a problem partly solved*, page 40–58. Cambridge University Press, 2006.

[112] V. Lucarini. Response operators for Markov processes in a finite state space: radius of convergence and link to the response theory for Axiom A systems. *Journal of Statistical Physics*, 162:312–333, 2016.

[113] V. Lucarini. Revising and extending the linear response theory for statistical mechanical systems: evaluating observables as predictors and predictands. *Journal of Statistical Physics*, 173:1698–1721, 2018.

[114] V. Lucarini, F. Ragone, and F. Lunkeit. Predicting climate change using response theory: global averages and spatial patterns. *Journal of Statistical Physics*, 166:1036–1064, 2017.

[115] N. J. Lutsko, I. M. Held, and P. Zurita-Gotor. Applying the Fluctuation–Dissipation Theorem to a two-layer model of quasigeostrophic turbulence. *Journal of Atmospheric Sciences*, 72:3161–3177, 2015.

[116] L. Magri, M. Bauerheim, F. Nicoud, and M. P. Juniper. Stability analysis of thermo-acoustic nonlinear eigenproblems in annular combustors. Part II. Uncertainty quantification. *Journal of Computational Physics*, 235:411–421, 2016.

[117] N. Makri. Information guided noise reduction for Monte Carlo integration of oscillatory functions. *Chemical Physical Letters*, 400:446–452, 2004.

[118] Y. Mao, G. Chen, and S. Lian. A novel fast image encryption scheme based on 3D chaotic Baker maps. *International Journal of Bifurcation and Chaos*, 14:3613–3624, 2004.

[119] S. A. Margulis. *Variational sensitivity analysis and data assimilation studies of the coupled land surface-atmospheric boundary layer system*. PhD thesis, Massachusetts Institute of Technology, 2002.

[120] R. May. Simple mathematical models with very complicated dynamics. *Nature*, 261:459–467, 1976.

[121] M. Mehta, A. K. Mittal, and S. Dwivedi. The double-cusp map for the forced Lorenz system. *International Journal of Bifurcation and Chaos*, 13:3029–3035, 2003.

[122] C. A. Morales, M. J. Pacifico, and E. R. Pujals. Singular hyperbolic systems. *Proceedings of the American Mathematical Society*, 127:3393–3401, 1999.

[123] G. Nagy. *Radon-Nikodym theorems*, pages 300–321. Kansas State University, 2021.

[124] C. Nation and D. Porras. Quantum chaotic fluctuation-dissipation theorem: Effective Brownian motion in closed quantum systems. *Physical Review E*, 99(052139), 2019.

[125] H. Natiq, M. R. K. Ariffin, M. R. M. Said, and S. Banerjee. Enhancing the sensitivity of a chaos sensor for Internet of things. *Internet of Things*, 7(100083), 2019.

[126] A. Ni. Hyperbolicity, shadowing directions and sensitivity analysis of a turbulent three-dimensional flow. *Journal of Fluid Mechanics*, 863:644–669, 2019.

[127] A. Ni. Fast linear response algorithm for differentiating stationary measures of chaos. *arXiv e-prints*, arXiv:2009.00595, 2021.

[128] A. Ni. Approximating linear response by nonintrusive shadowing algorithms. *SIAM Journal on Numerical Analysis*, 59:2843–2865, 2022.

[129] A. Ni and Q. Wang. Sensitivity analysis on chaotic dynamical systems by nonintrusive least squares shadowing (NILSS). *Journal of Computational Physics*, 347:56–77, 2017.

[130] E. J. Nielsen and W. K. Anderson. Aerodynamic design optimization on unstructured meshes using the navier-stokes equations. *AIAA Journal*, 37:1411–1419, 1999.

[131] N. Nikitin. Characteristics of the leading Lyapunov vector in a turbulent channel flow. *Journal of Fluid Mechanics*, 849:942–967, 2018.

[132] N. Oberoi. *Towards parametric sensitivity of high-fidelity turbulent simulations*. PhD thesis, University of Maryland, 2022.

[133] N. Oberoi and J. Larsson. pyChan: a Python code for simple channel flow direct numerical simulations (DNS) . URL: https://larsson.umd.edu/code/. Accessed: January 4, 2023.

[134] S. Olver. *Numerical approximation of highly oscillatory integrals*. PhD thesis, University of Cambridge, 2007.

[135] D. Pazó, I. G. Szendro, J. M. López, and M. A. Rodríguez. Structure of characteristic lyapunov vectors in spatiotemporal chaos. *Physical Review E*, 78:016209, 2008.

[136] C.-K. Peng, S. V. Buldyrev, S. Havlin, M. Simons, H. E. Stanley, and A. L. Goldberger. Mosaic organization of dna nucleotides. *Physical Review E*, 49:1685–1689, 1994.

[137] J. E. V. Peter and R. P. Dwight. Numerical sensitivity analysis for aerodynamic optimization: A survey of approaches. *Computers & Fluids*, 39:373–391, 2010.

[138] F. Pichler and J. Scharinger. Finite dimensional generalized baker dynamical systems for cryptographic applications. In F. Pichler, R. M. Díaz, and R. Albrecht, editors, *Computer Aided Systems Theory — EUROCAST '95*, pages 465–476, Berlin, Heidelberg, 1996. Springer Berlin Heidelberg.

[139] S. Y. Pilyugin. *Shadowing in dynamical systems*. Springer-Verlag New York, 1999. Volume 1706 of Lecture Notes in Mathematics.

[140] S. B. Pope. *Turbulent Flows*, pages 264–332. Cambridge University Press, Cambridge, UK, 2000.

[141] R. Pulch, E. J. W. ter Maten, and F. Augustin. Sensitivity analysis and model order reduction for random linear dynamical systems. *Mathematics and Computers in Simulations*, 111:80–95, 2015.

[142] C. Rackauckas, Y. Ma, V. Dixit, X. Guo, M. Innes, J. Revels, J. Nyberg, and V. Ivaturi. A comparison of automatic differentiation and continuous sensitivity analysis for derivatives of differential equation solution. *arXiv e-prints*, arXiv:1812.01892, 2018.

[143] F. Ragone, V. Lucarini, and F. Lunkeit. A new framework for climate sensitivity and prediction: a modelling perspective. *Climate Dynamics*, 46:1459–1471, 2016.

[144] C. H. Reick. Linear response of the Lorenz system. *Physical Review E*, 66(036103), 2002.

[145] L. Ren and M. Hartnett. Sensitivity analysis of a data assimilation technique for hindcasting and forecasting hydrodynamics of a complex coastal water body. *Computers & Geosciences*, 99:81–90, 2017.

[146] J. Reneaux and Thibert J.-J. The use of numerical optimization for airfoil design. *AIAA Journal*, 1985.

[147] D. Ruelle. A measure associated with Axiom-A attractors. *American Journal of Mathematics*, 98(3), 1976.

[148] D. Ruelle. Differentiation of SRB states. *Communications in Mathematical Physics*, 187:227–241, 1997.

[149] D. Ruelle. General linear response formula in statistical mechanics, and the fluctuation-dissipation theorem far from equilibrium. *Physics Letters A*, 245:220–224, 1998.

[150] D. Ruelle. Differentiation of SRB states: correction and complements. *Communications in Mathematical Physics*, 234:185–190, 2003.

[151] D. Ruelle. Differentiation of SRB states for hyperbolic flows. *Ergodic Theory and Dynamical Systems*, 28:613–631, 2008.

[152] D. Ruelle. A review of linear response theory for general differentiable dynamical systems. *Nonlinearity*, 22(2):855, 2009.

[153] D. Ruelle. Linear response theory for diffeomorphisms with tangencies of stable and unstable manifolds — a contribution to the gallavotti–cohen chaotic hypothesis. *Nonlinearity*, 31(12):5683–5691, 2018.

[154] K. Shawki and G. Papadakis. A preconditioned Multiple Shooting Shadowing algorithm for the sensitivity analysis of chaotic systems. *Journal of Computational Physics*, 398(108861), 2019.

[155] C. Shi-qing, Z. Shen-zhong, H. Yan-zhang, and Z. Wei-yao. Sensitivity coefficients of single-phase flow in low-permeability heterogeneous reservoirs. *Applied Mathematics and Mechanics (English Edition)*, 23:712–720, 2002.

[156] G. I. Sivashinsky. Nonlinear analysis of hydrodynamic instability in laminar flames — Part I. Derivation of basic equations. *Acta Astronautica*, 4:1177–1206, 1977.

[157] R. D. Skeel. What makes molecular dynamics work? *SIAM Journal on Scientific Computing*, 31:1363–1378, 2009.

[158] A. A. Śliwiak, N. Chandramoorthy, and Q. Wang. Ergodic Sensitivity Analysis of One-Dimensional Chaotic Maps. *Theoretical and Applied Mechanics Letters*, 10(6):438–447, 2020.

[159] A. A. Śliwiak, N. Chandramoorthy, and Q. Wang. Computational assessment of smooth and rough parameter dependence of statistics in chaotic dynamical systems. *Communications in Nonlinear Science and Numerical Simulation*, 101(105906), 2021.

[160] A. A. Śliwiak and Q. Wang. Differentiating densities on smooth manifolds. *Applied Mathematics and Computation*, 410(126444), 2021.

[161] A. A. Śliwiak and Q. Wang. Space-split algorithm for sensitivity analysis of discrete chaotic systems with unstable manifolds of arbitrary dimension. *arXiv e-prints*, arXiv:2109.13313, 2021.

[162] A. A. Śliwiak and Q. Wang. A trajectory-driven algorithm for differentiating SRB measures on unstable manifolds. *SIAM Journal on Scientific Computing*, 44(1):A312–A336, 2022.

[163] A. A. Śliwiak and Q. Wang. Approximating the linear response of physical chaos. *Nonlinear Dynamics*, 111:1835–1869, 2023.

[164] S. Smale. Differentiable dynamical systems. *Bulletin of the American Mathematical Society*, 73:747–817, 1967.

[165] G. Söderlind. Automatic control and adaptive time-stepping. *Numerical Algorithms*, 31:281–310, 2002.

[166] C. Sparrow. *The Lorenz Equations*. Springer-Verlag New York, 1982.

[167] N.-Z. Sun and W. W.-G. Yeh. Coupled inverse problems in groundwater modeling: 1. Sensitivity analysis and parameter identification. *Water Resources Research*, 26:2507–2525, 1990.

[168] I. Sunseri, J. Hart, B. Van Bloemen Waanders, and A. Alexanderian. Hyper-differential sensitivity analysis for inverse problems constrained by partial differential equations. *Inverse Problems*, 36(112), 2020.

[169] A. T. Takeuchi, H. Yang, F. Ginelli, G. Radons, and H. Chaté. Hyperbolic decoupling of tangent space and effective dimension of dissipative systems. *Physical Review E*, 84:046214, 2011.

[170] A. Tantet, V. Lucarini, F. Lunkeit, and H. A. Dijkstra. Crisis of the chaotic attractor of a climate model: a transfer operator approach. *Nonlinearity*, 31(2221), 2018.

[171] R. A. Thiétart and B. Forgues. Chaos theory and organization. *Organization Science*, 6(1):19–31, 1995.

[172] B. S. Thomson. Chapter 5 - differentiation. In E. PAP, editor, *Handbook of Measure Theory*, pages 179–247. North-Holland, Amsterdam, 2002.

[173] J. Thuburn. Climate sensitivities via a Fokker–Planck adjoint approach. *Quarterly Journal of the Royal Meteorological Society*, 131:73–92, 2005.

[174] S. Ulam. *Problems in Modern Mathematics*. Science Editions John Wiley & Sons, Inc., New York, 1964.

[175] R. Uma and D. Subbarao. Intrinsic chaos in a plasma wave. *Physics of Fluids B: Plasma Physics*, 2(1154), 1990.

[176] D. L. Van Kekem. Dynamics of the Lorenz-96 model: Bifurcations, symmetries and waves. *University of Groningen Research Database*, 2018.

[177] J. W. Vaupel and A. I. Yashin. Heterogeneity's ruses: some surprising effects of selection on population dynamics. *The American Statistician*, 39(3):176–185, 1985.

[178] H.X. Vo and L.J. Durlofsky. Data assimilation and uncertainty assessment for complex geological models using a new PCA-based parameterization. *Computational Geosciences*, 19:747–767, 2015.

[179] S. Wan, R. C. Sinclair, and P. V. Coveney. Uncertainty quantification in classical molecular dynamics. *Philosophical Transactions of the Royal Society A*, 379(20200082), 2021.

[180] Q. Wang. *Uncertainty quantification for unsteady fluid flow using adjoint-based approaches*. PhD thesis, Stanford University, 2009.

[181] Q. Wang. Convergence of the least squares shadowing method for computing derivative of ergodic averages. *SIAM Journal of Numerical Analysis*, 52:156–170, 2014.

[182] Q. Wang, R. Hu, and P. Blonigan. Least Squares Shadowing sensitivity analysis of chaotic limit cycle oscillations. *Journal of Computational Physics*, 267:210–224, 2014.

[183] X. Wang, M. Yoo, R. Glardon, and J. Furbringer. Performance and sensitivity analysis of supply chain alternative configurations - A case study. *2009 International Conference on Computers and Industrial Engineering*, pages 708–713, 2009.

[184] Z. Wei, I. Moroz, J. C. Sprott, A. Akgul, and W. Zhang. Hidden hyperchaos and electronic circuit application in a 5D self-exciting homopolar disc dynamo. *Chaos*, 27(033101), 2017.

[185] Z. Weimin, H. Jun, S. Jingsong, S. Jun, and S. Desheng. Improving product develop process time based on process sensitivity analysis. In *2012 IEEE International Conference on Computer Science and Automation Engineering (CSAE)*, volume 3, pages 76–79, 2012.

[186] R. F. Williams. Expanding attractors. *Publications Mathématiques de l'IHÉS*, 43:169–203, 1974.

[187] C. L. Wormell. Non-hyperbolicity at large scales of a high-dimensional chaotic system. *Proceedings of the Royal Society A*, 478, 2022.

[188] C. L. Wormell and G. A. Gottwald. On the validity of linear response theory in high-dimensional deterministic dynamical systems. *Journal of Statistical Physics*, 172:1479–1498, 2018.

[189] C. L. Wormell and G. A. Gottwald. Linear response for macroscopic observables in high-dimensional systems. *Chaos*, 29:21–30, 2019.

[190] M. Xu and M. R. Paul. Covariant lyapunov vectors of chaotic rayleigh-bénard convection. *Physical Review E*, 93:062208, 2016.

[191] X. I. A. Yang, J. Hong, M. Lee, and X. L. D. Huang. Grid resolution requirement for resolving rare and high intensity wall-shear stress events in direct numerical simulations. *Physical Review Fluids*, 6(054603), 2021.

[192] P. K. Yeung, K. R. Sreenivasan, and S. B. Pope. Effects of finite spatial and temporal resolution in direct numerical simulations of incompressible isotropic turbulence. *Physical Review Fluids*, 3(064603), 2018.

[193] L.-S. Young. Some large deviation results for dynamical systems. *Transactions of the American Mathematical Society*, 318(2):525–543, 1990.

[194] L.-S. Young. Statistical properties of dynamical systems with some hyperbolicity. *Annals of Mathematics*, 147:585–650, 1998.

[195] L.-S. Young. What Are SRB Measures, and Which Dynamical Systems Have Them? *Journal of Statistical Physics*, 108:733–754, 2002.

[196] D. Zhang, Y. Jiang, and J. Cai. Analytic Sensitivity Analysis for Shape Optimization. *Applied Mathematics and Mechanics (English Edition)*, 22:1325–1332, 2001.

[197] H. Zhang and A. Sandu. FATODE: A library for forward, adjoint, and tangent linear integration of ODEs. *SIAM Journal on Scientific Computing*, 36:C504–C523, 2014.