SVD algorithms to approximate spectra of dynamical systems ⋆

L. Dieci

School of Mathematics, Georgia Tech, Atlanta, GA 30332 U.S.A.

C. Elia

Dipartimento di Matematica, University of Bari, I-70100, Bari, Italy

Abstract

In this work we consider algorithms based on the Singular Value Decomposition (SVD) to approximate Lyapunov and Exponential Dichotomy spectra of dynamical systems. We review existing contributions, and propose new algorithms of the continuous SVD method. We present implementation details for the continuous SVD method, and illustrate on several examples the behavior of continuous (and also discrete) SVD method. This paper is the companion paper of [14].

Key words: Lyapunov exponents, exponential dicothomy, singular value decomposition, numerical approximation
1991 MSC: 15A18, 34D08, 65L07

1 Introduction

Investigation of a dynamical system, discrete or continuous, linear or nonlinear, deterministic or random, quite often requires information on the spectrum of the system. In this work, we will be specifically concerned with continuous deterministic dynamical systems:

$$\dot{x} = f(x) \ , \qquad t \geq 0 \ , \qquad x(0) = x_0 \ , \tag{1}$$

⋆ Work supported under NSF Grant DMS/FRG-0139895
Email addresses: dieci@math.gatech.edu (L. Dieci), elia@dm.uniba.it (C. Elia).

where $f : \mathbb{R}^n \to \mathbb{R}^n$ is Lipschitz and continuously differentiable, and we let $\phi^t(x_0)$ be the solution of (1), which we will assume to describe a bounded trajectory. Some powerful tools to analyze the nonlinear system (1) rely on the spectral information associated to the linearized system along $\phi^t(x_0)$:

$$\dot{\Phi}(t, x_0) = Df(\phi^t(x_0))\Phi(t, x_0) \ , \ \ \Phi(0, x_0) = I \ . \tag{2}$$

For the sake of generality, we will henceforth consider the linear non-autonomous system

$$\dot{X} = A(t)X, \qquad t \geq 0, \tag{3}$$

where $A : \mathbb{R}^+ \to \mathbb{R}^{n\times n}$ is continuous and bounded, and X is some fundamental matrix solution (i.e., $X(0)$ is invertible). For later reference, we will reserve the notation Φ to indicate the principal matrix solution of (3): $\Phi(0) = I$.

Our concern in this work will be to approximate the Lyapunov and the Exponential Dichotomy (or Sacker-Sell) spectra, Σ_{L} and Σ_{ED} respectively, of (3). These spectra will be defined shortly. For the moment, let us recall that these spectra play a fundamental role in classical stability studies (e.g., see [24], [12]); for example, study of variation with respect to the initial conditions for the nonlinear system $\dot{x} = f(x)$ leads naturally to study the Lyapunov spectrum, whereas study of variation with respect to parameters for $\dot{x} = f(\lambda, x)$ leads more naturally to the Exponential Dichotomy spectrum. Remarkably, and at first glance surprisingly, these well understood asymptotic features of a system reflect in many other contexts and Σ_{L} and Σ_{ED} are of use in perturbation and continuation of invariant manifolds (e.g.,see [46]), in shadowing and other error statements associated to numerical integration (see [38]), and of course spectral information is routinely used as indicator of chaotic behavior (see [36] for a very nice and accessible discussion of various issues on this), to extract entropy and other statistics of a dynamical systems (again, see [36]), to estimate dimension of attractors (e.g., [11]), as well as a tool for studying bifurcation theory of random dynamical systems [3], and of nonautonomous systems [20,27].

Many studies where spectra are used are computational in nature. This is unavoidable, of course, since differential equations cannot be solved exactly in the first place. In particular, some researchers ([23,22,43]) used an approach for approximating Lyapunov exponents based on the Singular Value Decomposition (SVD) of the linearized flow (i.e., of the fundamental matrix solution of the linearized problem). Albeit not as popular as methods based on the QR factorization of the matrix solution, the SVD approach presents some distinct potential advantages, as we will review in the course of this work.

The SVD approach has been traditionally motivated by the Multiplicative Ergodic Theorem (MET) of Oseledec, [35], and see Section 1.1.2 below. In [14], instead (and see also [16]), we have justified SVD methods based upon the

integral separation of the underlying linear system: This helps understanding how SVD methods will work in practice, and forms the backing for our algorithmic development. Still, regardless of the point of view one adopts to justify SVD techniques, we must warn that it is not easy at all to implement an SVD method. Our main purpose in this work is precisely that of proposing new algorithms for SVD methods, and give implementation details and show performance on a few examples.

A plan of this paper is as follows. In the remainder of this introduction, we review the concepts of Lyapunov and Exponential Dichotomy spectra, Σ_{L} and Σ_{ED}. We will also review an alternative to SVD methods for approximation of spectra, namely QR techniques. For reasons of brevity, we will give only a very brisk review, and refer to our companion paper [14] and the cited literature for details and theoretical justifications. In Section 2 we discuss the basic continuous and discrete SVD methods and their strengths and limitations. In Section 3 we present our new continuous SVD algorithm, and give implementation details. Finally, in Section 4 we show performance of SVD techniques on several linear and nonlinear examples.

1.1 Theoretical Background

To define Σ_{L}, we follow the classical definitions of Lyapunov, see [33]. Consider (3), and define the following quantities

$$\lambda_j^s = \limsup_{t\to+\infty} \frac{1}{t} \log \|X(t)e_j\| \ , \qquad \lambda_j^i = \liminf_{t\to+\infty} \frac{1}{t} \log \|X(t)e_j\| \ . \tag{4}$$

Here, and everywhere else, the vector norm is always the 2-norm, and the matrix norm is the induced norm; the vectors e_j, $j = 1, \ldots, n$, are the standard unit vectors. When $\sum_{j=1}^n \lambda_j^s$ (respectively, $\sum_{j=1}^n \lambda_j^i$) is minimized with respect to all fundamental matrix solutions of the system, the λ_j^s (resp., λ_j^i) are called *upper* (*lower*) *Lyapunov exponents* (LEs for short) and the corresponding X (i.e., $X(0)$) is said to form a *normal basis*. [The existence of a normal basis was established by Lyapunov himself, see [33]. Presently, we observe that there are several normal bases.] For so-called *regular* systems, one has $\lambda_j^s = \lambda_j^i$, for all $j = 1, \ldots, n$. Naturally, we can rephrase (4) in terms of the principal matrix solution. That is, if $X(0)$ is a normal basis, letting $\overline{x}_j$, $j = 1, \ldots, n$, be its columns, then the upper/lower Lyapunov exponents are given by

$$\lambda_j^s = \limsup_{t\to+\infty} \frac{1}{t} \log \|\Phi(t)\overline{x}_j\| \ , \qquad \lambda_j^i = \liminf_{t\to+\infty} \frac{1}{t} \log \|\Phi(t)\overline{x}_j\| \ . \tag{5}$$

The Lyapunov spectrum is defined to be

$$\Sigma_{\mathrm{L}} = \bigcup_{j=1}^{n} [\lambda_j^i, \lambda_j^s] \ . \tag{6}$$

The Exponential Dichotomy, or Sacker-Sell, spectrum of (3), denote it with Σ_{ED}, is the set of all real values λ such that the shifted system

$$\dot{x} = (A(t) - \lambda I)x, \tag{7}$$

does not admit exponential dichotomy (see [41]). Recall that (3) admits an exponential dichotomy in $[0, +\infty)$ if there exists a projection P and real numbers $K \geq 1$, $\alpha > 0$ s.t. for any fundamental matrix solution X we have

$$\begin{aligned} \|X(t)PX(s)^{-1}\| &\leq \ Ke^{-\alpha(t-s)}, \qquad \forall t, s, \quad t \geq s \geq 0 \ , \\ \|X(t)(I-P)X(s)^{-1}\| &\leq \ Ke^{\alpha(t-s)}, \qquad \forall t, s, \quad 0 \leq t \leq s \ . \end{aligned} \tag{8}$$

Σ_{ED} is the union of m disjoint closed intervals, $m \leq n$:

$$\Sigma_{\mathrm{ED}} = \bigcup_{i=1}^{m} [a_i, b_i] \ : \quad a_1 \leq b_1 < a_2 \leq b_2 < \cdots < a_m \leq b_m \, . \tag{9}$$

In general, Σ_{ED} is always stable with respect to perturbations, by virtue of the "Roughness Theorem" for Exponential Dichotomies (see [12]). Σ_{L}, instead, enjoys only conditional stability. Necessary and sufficient conditions for Σ_{L} to be stable are given in [2]. In case in which the Lyapunov exponents are distinct, the necessary and sufficient condition for stability of Σ_{L} is that (3) be integrally separated. This means that there exists a fundamental matrix solution X of (3) and constants $a > 0$ and $0 < d \leq 1$ such that for all $i = 1, \ldots, n-1$, we have

$$\frac{\|X(t)e_i\|}{\|X(s)e_i\|} \frac{\|X(s)e_{i+1}\|}{\|X(t)e_{i+1}\|} \geq de^{a(t-s)}, \qquad \forall t, s, \quad t \geq s \geq 0 \ . \tag{10}$$

1.1.1 Directions: Growth Subspaces

Just like eigenvalues of a matrix have geometrical informations associated to them, the eigendirections, the growth factors in Σ_{L} and Σ_{ED} can be supplemented by geometrical information on the subspaces of solutions which achieve a specific growth. This is an important point with far reaching ramifications. Consider first Σ_{L}, and the upper Lyapunov exponents. Let λ_j^s, $j = 1, \ldots, p$, be the distinct upper LEs of (3), of multiplicities $d_j, j = 1, \ldots, p$. For $j = 1, \ldots, p$, define W_j to be the set of all initial conditions w such that for the solution $\Phi(t)w$, we have $\limsup_{t\to\infty} \frac{1}{t} \log \|\Phi(t)w\| \leq \lambda_j^s$. These W_j's are $(d_1 + \cdots + d_j)$-dimensional linear spaces corresponding to the max number of linearly inde-

pendent solutions with exponent equal to λ_j^s. The W_j's form a filtration of $\mathbb{R}^n$: $\mathbb{R}^n = W_1 \supset W_2 \ldots \supset W_p \supset W_{p+1} = \{0\}$, and

$$y(t) \in W_j \backslash W_{j+1} \Longleftrightarrow \limsup_{t\to\infty} \frac{1}{t} \log(\|y(t)\|) = \lambda_j^s \,. \tag{11}$$

Alternatively, if we let V_j be the orthogonal complement of W_{j+1} in W_j, i.e. $W_j = W_{j+1} \oplus V_j$, $V_j \perp W_{j+1}$, then $\mathbb{R}^n = V_1 \oplus V_2 \oplus \ldots \oplus V_p$, and if $w \in V_j$ then $\limsup_{t\to\infty} \frac{1}{t} \log \|\Phi(t)w\| = \lambda_j^s$. Notice that if we have distinct LEs, then $\dim(V_j) = 1$ for all $j = 1, \ldots, n$, and we will simply write $\overline{v}_j$ for the normalized directions in this case. In [14], we have shown that the $\overline{v}_j$'s, for integrally separated systems, characterize the set of initial conditions leading to lower Lyapunov exponents as well.

Also to each interval of $\Sigma_{\rm ED}$ one can associate linear spaces $\mathcal{N}_j$, $j = 1 \ldots, m$; see [41]. In [14] we proved that these are the subspaces of initial conditions giving growth behavior between a_j and b_j of (9), uniformly. That is, $w \in \mathcal{N}_j$, $w \neq 0$, if and only if for some constants $K_j \geq 1$ and $K_{j-1} \geq 1$ we have

$$\frac{\|\Phi(t)w\|}{\|\Phi(s)w\|} \leq K_j e^{b_j(t-s)} \;, \;\; \text{and} \;\; \frac{1}{K_{j-1}} e^{a_j(t-s)} \leq \frac{\|\Phi^{-T}(s)w\|}{\|\Phi^{-T}(t)w\|}$$

for all t, s: $t \geq s \geq 0$. For integrally separated systems, in [14] we proved that the $\mathcal{N}_j$'s are spanned by the vectors $\overline{v}_k$'s associated to the Lyapunov exponents inside the interval $[a_j, b_j]$.

1.1.2 Nonlinear case

When dealing with the nonlinear system (1), the spectra must be understood as being associated to the linearized problem (2). Therefore, in general, they will depend on the initial condition x_0. However, many nonlinear systems of practical interest have solutions which eventually settle on an attractor. In this case, important measure theoretic statements can be made for the spectra. The MET [1], says that if the flow of (1) is taking place on a compact manifold and μ is an ergodic invariant measure for the manifold, then there is a set B of full measure such that if x_0 belongs to B, then the Lyapunov exponents become independent of x_0, and further exist as limits. Similarly, $\Sigma_{\rm ED}$ becomes independent of x_0. In such cases, one can properly speak of $\Sigma_{\rm L}$ and $\Sigma_{\rm ED}$ for the nonlinear system (1). We recall that an important consequence of the MET is that the Lyapunov exponents are given by the logarithms of the eigenvalues of the following limit matrix $\Lambda(x_0)$ (this limit exists for almost any initial condition x_0 with respect to an invariant measure, and further the limit $\Lambda(x_0)$

[1] The largest part of [3] is devoted to this remarkable theorem.

becomes independent of x_0 when the measure is ergodic):

$$\lim_{t\to\infty}(\Phi(t,x_0)^T\Phi(t,x_0))^{1/2t} = \Lambda(x_0)\ . \tag{12}$$

1.2 Computational Background

The outstanding difficulty in approximating spectra is to avoid direct integration of fundamental matrix solutions, since solution modes tend to align in the direction of fastest growth. Any successful method will need to avoid this phenomenon.

Two families of techniques have been explored, based on the QR and SVD decompositions of the fundamental matrix solution. The general idea is to transform the problem to one with a simpler structure, triangular or diagonal, and to extract the spectral information from the diagonal subsystems. For numerical reasons one favors use of orthogonal transformations. For simplicity, we describe the case in which we want to approximate all spectral intervals. In Section 2.3, we outline the modifications which are needed in case we want to approximate only a few dominant spectral intervals. Below, we describe QR methods, in the next Chapter we describe the SVD methods.

For a fundamental matrix solution X, the idea is to obtain the decomposition $X = QR$ with Q orthogonal and R upper triangular with positive diagonal: $R_{ii} > 0$. There are two classes of methods: *discrete* and *continuous* QR methods. Both are widely used, especially the discrete method, and there are even public domain codes based on these methods; e.g., see [6,10,16,22] and the bibliography in the technical report [19], which also serves as a user's guide for available public domain codes.

1.2.1 Discrete QR

Let $t_0 = 0$, $X_0 = Q_0R_0$, and let $t_0 < t_1 < t_2 < \ldots$, be the points of the computational grid (it could be an equispaced grid, $t_k = kh$, for $k = 0, 1, \ldots$, for given $h > 0$, but more typically uneven grids are used).

For a general grid point t_{k+1}, we want the QR factorization of X at t_{k+1}. In terms of transition matrices, we can write

$$X(t_{k+1}) = \Phi(t_{k+1},t_k)\cdots\Phi(t_2,t_1)\Phi(t_1,t_0)Q_0R_0\,.$$

Recall that, for $j = 0, 1, \ldots$, $\Phi(t_{j+1},t_j)$ is the solution at t_{j+1} of the problem

$$\frac{d}{dt}\Phi(t,t_j) = A(t)\Phi(t,t_j)\ , \quad \Phi(t_j,t_j) = I\,.$$

Thus, to obtain the QR factorization of $X(t_{k+1})$ one can proceed as follows

$$\Phi(t_1, t_0)Q_0 = Q_1R_1\,, \quad \Phi(t_2, t_1)Q_1 = Q_2R_2\,, \quad \dots\,,$$

so that

$$X(t_{k+1}) = Q_{k+1}R_{k+1}\dots R_1R_0\,. \tag{13}$$

Approximation to the (upper) Lyapunov exponents can be obtained from

$$\lambda_i^s = \limsup_{t_k\to\infty} \frac{1}{t_k}\sum_{j=0}^{k} \log(R_j)_{ii}\,, i = 1, \dots, n\,.$$

Approximation of the endpoints of the spectral intervals of Σ_{ED} can also be based solely on the diagonal of R; for details, see [19] and [16].

Remark 1.1 Observe that it is easy to proceed from the QR factorization of X at t_{k+1} to the QR factorization at t_{k+2}, and so forth.

1.2.2 Continuous QR

Let $t_0 = 0$, $X_0 = Q_0R_0$. The method consists in defining and integrating a differential equation for Q. Differentiating the relation $X = QR$, one gets $\dot{Q}R + Q\dot{R} = A(t)QR$. Multiplying this by Q^T on the left gives

$$\dot{R} \;=\; B(t)R\,, \quad R(0) = R_0\,, \;\text{ where }\; B(t) := Q^TA(t)Q - S\,,$$

and we have formally set $S := Q^T\dot{Q}$. Similarly,

$$\dot{Q} \;=\; AQ - QB = QS\,, \quad Q(0) = Q_0\,.$$

Observe that S is skew-symmetric and B must be upper triangular, and so

$$S_{ij} = \begin{cases} (Q^T(t)A(t)Q(t))_{ij}, & i > j, \\ 0, & i = j, \\ -S_{ji}, & i < j. \end{cases}$$

The upper LEs can be obtained from

$$\begin{aligned} \lambda_i^s &= \limsup_{t\to\infty} \frac{1}{t}\int_0^t (Q^T(s)A(s)Q(s))_{ii}ds \\ &= \limsup_{t\to\infty} \frac{1}{t}\int_0^t B_{ii}(s)ds\,, \qquad i = 1, \dots, n\,. \end{aligned}$$

Remark 1.2 Needless to say, one needs to integrate all relevant differential equations numerically, and there are several subtleties involved; e.g., to maintain orthogonality of the computed Q in the continuous QR method. For details about state of the art implementations of discrete and continuous QR methods, we again refer to the report [19].

2 SVD methods

Even here there are discrete and continuous methods. The discrete SVD method is based on the so–called *product SVD algorithm*, which has found applications also in contexts different than approximation of Lyapunov exponents, and it has thus attracted the attention of several researchers in computational linear algebra; to witness, while the works [1,34,43] investigate discrete SVD methods for the purpose of approximation of Lyapunov exponents, the works [7,8,25,44] study discrete SVD methods for completely different applications. We would also like to draw attention to the work of Lust, [32], who developed codes for the so-called periodic Schur decomposition, a close relative to the discrete SVD method, for extracting Floquet multipliers. In the context of approximation of Lyapunov exponents, algorithms based on the continuous SVD method have been much less explored; the work [23] is an early example, and our work [15] is a specific continuous SVD method for a structured problem. The only work of which we are aware which considers both continuous and discrete SVD methods for approximation of Lyapunov exponents of regular systems is the review [21]. No work of which we are aware considers SVD methods (discrete or continuous) for approximation of the $\liminf$'s and $\limsup$'s in Σ_{L}, nor of Σ_{ED}.

2.1 Discrete SVD

The setup is similar to that of the discrete QR method, see Section 1.2.1. We want the SVD factorization of $\Phi(t)$ at t_{k+1}. Since

$$\Phi(t_{k+1}) = \Phi(t_{k+1}, t_k) \cdots \Phi(t_2, t_1)\Phi(t_1, t_0)\,,$$

the task is how to compute the SVD of $\Phi(t_{k+1})$ without forming the product explicitly (both to avoid catastrophic overflows and numerical dependency of the columns). The basic tool is the so-called periodic (or product) Schur decomposition.

For simplicity, let the above product be rewritten as

$$\Phi^{(k+1)}\Phi^{(k)} \ldots \Phi^{(2)}\Phi^{(1)}\,.$$

The periodic Schur decomposition consists in applying the classical QR (or Schur) algorithm simultaneously to

$$\left(\Phi^{(k+1)}\Phi^{(k)} \ldots \Phi^{(2)}\Phi^{(1)}\right)^T \left(\Phi^{(k+1)}\Phi^{(k)} \ldots \Phi^{(2)}\Phi^{(1)}\right),$$

and to

$$\left(\Phi^{(k+1)}\Phi^{(k)} \ldots \Phi^{(2)}\Phi^{(1)}\right)\left(\Phi^{(k+1)}\Phi^{(k)} \ldots \Phi^{(2)}\Phi^{(1)}\right)^T,$$

by doing one QR factorization at the time. A synthetic description of this process, which avoids explicitly forming the above products, is given in [43] and goes as follows.

Algorithm 2.1 *Product SVD.*

- Initialize: $U_1 = I$, $V_1 = I$, $\Phi_1^{(p)} = \Phi^{(p)}$ for $p = 1, 2, \ldots, k+1$.

- (Iterate) For $i = 1, 2, \ldots,$

$$
\begin{aligned}
&Q_i^{(0)} \leftarrow I \\
&\text{For } p = 1, \ldots, k+1 \\
&\qquad C_i^{(p)} \leftarrow \Phi_i^{(p)} Q_i^{(p-1)} \ , \\
&\qquad C_i^{(p)} = Q_i^{(p)} R_i^{(p)} \ . \quad \text{[QR factorization]}
\end{aligned}
$$

- (Update) If i is odd, then $U_{i+1} \leftarrow U_i Q_i^{(k+1)}$, else $V_{i+1} \leftarrow V_i Q_i^{(k+1)}$
- (Reverse) For $p = 1, 2, \ldots, k+1$ replace $\Phi_{i+1}^{(p)} \leftarrow \left(R_i^{(k+1-p+1)}\right)^T$

Iterating the above process to the limit, as for the Schur algorithm, and calling $U^{(k+1)}$ and $V^{(k+1)}$ the obtained (approximations to the) limits of U_i and V_i, respectively, and $R^{(p)}$ the associated limiting triangular factors of $R_i^{(p)}$, $p = 1, 2, \ldots, k+1$, one eventually finds the product SVD factorization as

$$
U^{(k+1)} R^{(k+1)} R^{(k)} \ldots R^{(2)} R^{(1)} (V^{(k+1)})^T \ , \tag{14}
$$

where the product of the triangular factors is diagonal (cfr. with the discrete QR case, where the product of the triangular factors in (13) remains triangular). Of course, only the diagonal of the factors in (14) needs to be monitored for extracting spectral information. In particular, approximation to the (upper) Lyapunov exponents can be obtained from

$$
\lambda_i^s = \limsup_{t_k \to \infty} \frac{1}{t_k} \sum_{j=0}^{k} \log(R^{(j)})_{ii} \, , i = 1, \ldots, n \, .
$$

Remarks 2.2 Several observations are in order.

(i) It is not clear how to exploit knowledge of the SVD of $\Phi(t_{k+1})$ to obtain the SVD of Φ at a later time, say of $\Phi(t_{k+2})$. Or, better, one possible source of computational savings is in bypassing the first cycle of QR factorizations. In fact, in the initialization phase of Algorithm 2.1, we can set $V_1 = V^{(k+1)}$, $\Phi_1^{(p)} = R^{(p)}$, $p = 1, 2, \ldots, k+1$, and $\Phi_1^{(k+2)} = \Phi^{(k+2)}$. This will save one sweep of QR factorizations. No additional computational savings appear possible.

(ii) It is not easy to see how to avoid storing a lot of matrices, albeit by virtue of the above point (i), we can assume that they are all (but a few of them)

triangular. For a product of length N, we therefore need to store about $N(n^2+n)/2$ floating point numbers. As a consequence, the technique is not feasible for large problems and long products, because of high memory requirement.

(iii) On the surface, the method does not seem to require distinct singular values at any given time. However, this is misleading. Indeed, available theoretical justifications for algorithms to approximate spectra based on the SVD decomposition require this assumption (see [14]). Moreover, the finite precision error analysis of Oliveira & Stewart (see [34]) of the product SVD algorithm needs an even more stringent assumption.

(iv) In Section 4, Example 4.3, we approximate the directions leading to the different LEs via the discrete SVD. Our implementation is not particularly sophisticated, although shifting strategies to accelerate convergence of the periodic Schur decomposition have been used, but it is nevertheless indicative of the behavior of discrete SVD methods.

2.2 Continuous SVD

To avoid the drawbacks of the discrete SVD method, see especially Remarks 2.2-(i,ii), we have turned our attention to continuous SVD techniques whereby we seek a smooth SVD of Φ: $\Phi(t) = U(t)\Sigma(t)V^T(t)$, for all t.

In general, it is well understood that a smooth SVD will not exist, at least not with the same degree of smoothness of Φ, see [13, Theorem 3.6]. However, in [13, Theorem 4.3], the authors showed that generic one parameter functions (for us, fundamental solutions) have distinct singular values. Moreover, if the singular values are distinct, and thus ordered for all t: $\sigma_1(t) > \ldots > \sigma_n(t)$, then the SVD factors are as smooth as Φ itself, and one can write differential equations for U, Σ, V. Now, assuming that the singular values are distinct, the differential equations for the factors U, Σ, V have been derived many times before; e.g., see [13,14,16,21,23,47]. They read

$$\begin{aligned} &\dot{\sigma} = \mathrm{diag}(C)\sigma \ , \quad \sigma = \mathrm{diag}(\Sigma) \ , \quad C := U^T A U \ , \\ &\dot{U} = UH \ , \\ &\dot{V} = VK \ , \end{aligned} \tag{15}$$

where H and K are antisymmetric with entries (for $i, j = 1, \ldots, n$, $i \neq j$)

$$\begin{aligned} H_{ij} &= \frac{c_{ij}\sigma_j^2 + c_{ji}\sigma_i^2}{\sigma_j^2 - \sigma_i^2} \ , \\ K_{ij} &= \frac{\Big(c_{ij} + c_{ji}\Big)\sigma_i\sigma_j}{\sigma_j^2 - \sigma_i^2} \ . \end{aligned} \tag{16}$$

The solutions of these differential equations, for given initial conditions U_0, Σ_0, V_0, will provide the sought smooth SVD. To support these differential equations models, we also recall that in [14, Theorems 4.2 and 4.6] we proved the following results:

Facts 2.3 Let system (3) have stable and distinct Lyapunov exponents (equivalently, the original problem has the property of integral separation). Then: (1) the singular values of any fundamental matrix solution will be distinct for $t \geq \overline{t}$, where $\overline{t}$ is finite (its value depends on the fundamental matrix); (2) from the singular values one can obtain Σ_{L} and Σ_{ED}, in particular approximation to the (upper) Lyapunov exponents can be obtained from

$$\lambda_i^s = \limsup_{t\to\infty} \frac{1}{t} \log \sigma_i(t) \, , i = 1, \ldots, n \, ; \tag{17}$$

(3) finally, the orthogonal function V converges exponentially fast to a constant matrix $\overline{V}$ whose columns are the directions of Section 1.1.1. In [14, Theorem 5.4] we showed that the rate of exponential convergence is

$$\alpha = A \max_{1\leq i\leq n-1} (\lambda_{i+1}^s - \lambda_i^s) \, , \quad 0 < A \leq 1 \ , \tag{18}$$

and $A = 1$ for regular systems (see [14, Corollary 5.5]).

Remarks 2.4 The following observations are in order.

(i) A major benefit of the continuous SVD formulation, versus the discrete SVD approach, is that the memory requirements are minimal, and it is (in principle) straightforward to pass from the SVD at one point to the SVD at some later point: Just keep on integrating the differential equations. Moreover, we observe that if all one needs are the Lyapunov exponents, then we do not need to find V; see (15).

(ii) But, there are also some potential difficulties when adopting a continuous SVD method as in (15).
   (a) If σ_i and σ_j get very close to each other, then H and K will change rapidly, and U and V will become harder to find.
   (b) For long time integration, some of the σ_i's will most likely overflow, inhibiting accurate numerical integration of (15).

(c) In theory, even for systems with stable Lyapunov exponents, the σ_i's will be distinct only for sufficiently large time, that is for $t \geq \overline{t}$ of Facts 2.3.

In Section 3, we will discuss in which way we overcame these difficulties.

2.3 Few Intervals

In several circumstances, it will not be necessary to approximate the entire spectrum of the system, say all of its Lyapunov exponents, but only a few of the most dominant intervals will suffice. For example, to estimate the entropy of a system only the positive Lyapunov exponents are necessary, while to estimate the dimension of an attractor only those leading to a positive sum are needed; see [37]. Furthermore, many systems possess special symmetries which reflect in symmetries in the spectrum (either Σ_{L} or Σ_{ED}), thus making it possible to recover approximation of the entire spectrum by approximating only part of it. E.g., in the important case of a Hamiltonian system, we only need to approximate half of the spectrum, because of its symmetry with respect to the origin. In all such cases, the SVD approaches can be modified so to avoid finding the SVD of the entire matrix solution Φ.

So, let us suppose that we want to approximate only p, with $p < n$, most dominant spectral intervals of (3). Let us also assume that $X_0 \in \mathbb{R}^{n\times p}$ are initial conditions in the p most dominant directions of the spectral intervals. Then, we will consider the function $X(t) = \Phi(t)X_0 \in \mathbb{R}^{n\times p}$, for all t, and will be seeking the (reduced) SVD of X: $X(t) = U(t)\Sigma(t)V^T(t)$, where $U(t) \in \mathbb{R}^{n\times p}$, $U^T(t)U(t) = I_p$, $\Sigma(t) \in \mathbb{R}^{p\times p}$, $\Sigma(t) = \mathrm{diag}\Big(\sigma_1(t), \ldots, \sigma_p(t)\Big)$, and $V(t) \in \mathbb{R}^{p\times p}$, $V^T(t)V(t) = I_p$, for all t. From the singular values $\sigma_1, \ldots, \sigma_p$, we will be able to extract the desired dominant spectral information.

Remark 2.5 A natural mathematical question is how to select the initial conditions X_0 so that the growth factors (i.e., the spectral intervals) extracted from $X(t) = \Phi(t)X_0$ will indeed be the p most dominant ones. Of course, if we knew the matrix $\overline{V}$, limit of the functions $V(t)$ of which in Remark 2.4-(i), we could certainly choose $X_0 = \overline{V}_{1:p}$; i.e., the first p columns of $\overline{V}$ will do. Since we do not know $\overline{V}$, we actually choose p random vectors in $\mathbb{R}^n$. This is the accepted practice; see [6,21]. The reason for the practical success of this approach is that the probability that the columns of X_0 have no component in the directions of the first p most dominant directions is 0. [In practice, we may as well confide in numerical errors and just take $X_0 = \left(\begin{smallmatrix} I_p \\ 0 \end{smallmatrix}\right)$.]

We now discuss how the discrete and continuous SVD methods need to be modified to achieve the desired reduced SVD.

2.3.1 Discrete SVD: Reduced Case

We want the reduced SVD factorization of $X(t)$ at $t = t_{k+1}$, avoiding to work with square matrices in $\mathbb{R}^{n\times n}$. We can achieve this as follows. Observe that

$$X(t_{k+1}) = \Phi(t_{k+1}, t_k)\cdots\Phi(t_2, t_1)\Phi(t_1, t_0)X_0\,, \quad X_0 \in \mathbb{R}^{n\times p}\,,$$

and the task is how to compute the SVD of $X(t_{k+1})$ (which is in $\mathbb{R}^{n\times p}$) without forming products and without finding the transition matrices $\Phi(t_j, t_{j-1})$. The procedure we suggest below is very simple, and it is analogous to what one does with the discrete QR method.

Let $X_0 = Q_0R_0$, with $Q_0 \in \mathbb{R}^{n\times p}$, $Q_0^TQ_0 = I_p$ and R_0 upper triangular in $\mathbb{R}^{p\times p}$ with positive diagonal. Let $X(t, t_0) = \Phi(t, t_0)Q_0$, for $t \in [t_1, t_0]$. Then, recursively solve these problems:

For $j = 1, \ldots, k+1$, solve
$\partial_t X(t, t_{j-1}) = A(t)X(t, t_{j-1})\ ,\ \ t_{j-1} \le t \le t_j\ ,\ \ X(t_{j-1}, t_{j-1}) = Q_{j-1}\ ,$
let $X(t_j, t_{j-1}) = Q_jR_j\ ,$ End

where in the QR factorizations above we have $Q_j \in \mathbb{R}^{n\times p}$, $Q_j^TQ_j = I_p$, and $R_j \in \mathbb{R}^{p\times p}$, upper triangular with positive diagonal. At the end of this process, we obtain that

$$X(t_{k+1}) = Q_{k+1}R_{k+1}\cdots R_2R_1R_0\,,$$

and so we can focus on finding the SVD factorization of the product $R_{k+1}\cdots R_2R_1R_0$, where all matrices are now square (p, p), and triangular. The procedure now is identical to the general square case, see Algorithm 2.1. The memory requirements have been reduced, but they are still very demanding: For a product of length N, we need to store N triangular factors in $\mathbb{R}^{p\times p}$ for a total of $N(p^2 + p)/2$ floating point numbers.

2.3.2 Continuous SVD: Reduced Case

We derive differential equations for the reduced SVD. The derivation is straightforward. Recall that we have $\dot{X} = A(t)X$. From the relation $X = U\Sigma V^T$, differentiating both sides, we obtain $AU\Sigma V^T = \dot{U}\Sigma V^T + U\dot{\Sigma}V^T + U\Sigma\dot{V}^T$. Letting again $\sigma = \mathrm{diag}(\Sigma)$, $C = U^TAU$, and $K = V^T\dot{V}$, $H = U^T\dot{U}$, we obtain the following set of differential equations

$$\begin{aligned}
&\dot{\sigma} = \mathrm{diag}(C)\sigma\ ,\quad \sigma = \mathrm{diag}\big(\sigma_1, \ldots, \sigma_p\big)\\
&\dot{U} = UH + (I - UU^T)AU\ ,\\
&\dot{V} = VK\ ,
\end{aligned} \tag{19}$$

where H and K are antisymmetric with entries specified in (16) (for $i, j = 1, \ldots, p$, $i \neq j$).

Remark 2.6 An alternative to the differential equations (19) consists in the following two step strategy. [We have not implemented this strategy, but it may be an interesting approach since with this technique one can use tried and true codes for the continuous QR as a building block; see [19].]

(i) Find the continuous QR factorization of X: $X = QR$. Here, Q will satisfy $\dot{Q} = AQ - QB = AQ - Q(Q^T AQ - S)$, where B, S are defined as in Section 1.2.2, but are now $(p \times p)$ matrices.
(ii) Use the continuous SVD method on the new problem $\dot{R} = B(t)R$, which is now a square problem of reduced dimensionality. On this square problem one would proceed just as we did in Section 2.2, obtaining (15) with B replacing A there.

2.4 QR vs SVD

Consider again the square case, that is when we want all spectral intervals. As we have reviewed, there are two distinct classes of methods, based on the QR and the SVD of a matrix solution. We believe that QR and SVD methods have some distinct advantages with respect to each other, and we view these methods as complementing each other. In general terms, QR methods are superior in terms of overall simplicity of implementation and general justification (see [18]), as long as the fundamental matrix on which they are used is a normal matrix solution. SVD methods, instead, have the advantage of working, in principle, with any matrix solution, but are more complicated to implement, and presently fully justified only for integrally separated systems, in which case they also provide an orthogonal representation for the growth subspaces (the directions of Section 1.1.1).

3 A new continuous SVD algorithm

In this section we discuss new SVD algorithms for the approximation of stability spectra of system (3). We will take the point of view that the entire spectra are desired, though the techniques below are easily modified in case we only desire a few most dominant spectral intervals; see Section 2.3.

As already remarked, to obtain feasible results, we need the assumption of stable and distinct LEs. Hence, from now on, this will be assumed to be the case. In particular, all theoretical results of [14, Sections 4 and 5] apply (see Facts

2.3), and we may as well consider the SVD for the principal matrix solution Φ, since the spectra are independent of the matrix solution we consider.

3.1 A new continuous SVD formulation

We need to address the potential difficulties of the continuous SVD method (15) expressed in points (a,b,c) of Remarks 2.4-(iii). First, let us address points (a) and (c).

The equations (15) for the SVD require distinct singular values. Now, recall Facts 2.3, we know that this will be the case for $t \geq \overline{t}$, though of course we do not know $\overline{t}$ before hand for a specific problem. Still, suppose we are looking at $t \geq \overline{t}$. In this case, the singular values are more than merely distinct, they are exponentially separated. In fact, see [14], we will have that

$$\frac{\sigma_j(t)}{\sigma_{j+1}(t)} \geq e^{at}, \quad \forall j = 1, \ldots, n-1, \text{and } \forall t \geq \overline{t} \ ,$$

for some constant $a > 0$. Therefore, for $t \geq \overline{t}$, point (a) of Remarks 2.4-(iii) does not apply.

Consider now the point (c) of Remarks 2.4-(iii). The concern here is that the singular values may coincide for some $0 < t < \overline{t}$, since we know that they are distinct past $\overline{t}$. However, this also turns out to be not a true concern in practice. To appreciate why, we need to appeal to a genericity statement. That is, we consider generic [2] full rank functions (for us, fundamental matrix solutions). Then, from the proof of [13, Theorem 4.3], we know that functions with two equal singular values form a stratified manifold with strata of dimension less than or equal to n^2-2; thus, having two equal singular values is a co-dimension 2 phenomenon and generic one-parameter (here, t) functions do not have two equal singular values. In practice, we just integrate for the fundamental matrix solution Φ on a small interval, say we solve $\dot{\Phi} = A(t)\Phi$, up to $t = t_\epsilon$, compute the algebraic SVD of $\Phi(t_\epsilon)$ and initialize the factors for the continuous SVD method at t_ϵ. In all problems we have solved, this straightforward procedure was successful.

Remark 3.1 Of course, there are many problems with special symmetries (hence, not generic) for which singular values coalesce. If we are aware of these symmetries, then special SVD procedures may be possibly used to avoid the difficulties caused by coalescing singular values; e.g., see Example 4.5 below, and see [15].

[2] with respect to the Whitney topology

Now let us turn our attention to the most serious difficulty, point (b) of Remarks 2.4-(iii). This is a true concern, since naive implementation of (15) does indeed lead to overflow problems very quickly. To overcome this issue, we change variables. Rather than working with the singular values σ_j's, we work with the variables

$$\nu_j(t) = \frac{\sigma_{j+1}(t)}{\sigma_j(t)}, \quad j = 1, \dots, n-1, \quad \nu_n(t) = \log(\sigma_n(t)) \ . \tag{20}$$

Observe that, being the singular values arranged in decreasing order, for $j = 1, \dots, n-1$, we have

$$0 < \nu_j(t) < 1 \ .$$

The differential equations for the ν_j's become

$$\begin{aligned} \dot{\nu}_j &= (c_{j+1,j+1} - c_{jj})\nu_j, \quad j = 1, \dots, n-1 \ , \\ \dot{\nu}_n &= c_{nn} \ . \end{aligned} \tag{21}$$

As it turns out, see Section 3.1.1 below, it will be even more convenient for us to work with the logarithms of the ν_j's, $j = 1, \dots, n-1$. That is, we will integrate the equations

$$\begin{aligned} \frac{d}{dt}(\log \nu_j) &= c_{j+1,j+1} - c_{jj}, \qquad j = 1, \dots, n-1, \\ \dot{\nu}_n &= c_{nn} \ . \end{aligned} \tag{22}$$

Finally, we need to rewrite equations (16) in terms of the variables ν_j's. To do so, define the following quantities, for $j > i$:

$$\nu_{ij}(t) = \frac{\sigma_j(t)}{\sigma_i(t)} = \prod_{k=j-1}^{i} \nu_k(t).$$

Note that these can be evaluated directly from the ν_i's. The equations for the entries of H and K are now rewritten as

$$\begin{aligned} H_{ij} &= \frac{c_{ij}\nu_{ij} + c_{ji}}{\nu_{ij}^2 - 1} \ , \quad j > i \ , \qquad H_{ij} = -H_{ji} \ , \quad j < i \\ K_{ij} &= \frac{c_{ij} + c_{ji}}{\nu_{ij}^2 - 1}\nu_{ij} \ , \quad j > i \ , \quad K_{ij} = -K_{ji} \ , \quad j < i \ . \end{aligned} \tag{23}$$

We finally summarize the skeleton of our continuous SVD algorithm:

- Integrate for Φ up to time t_ϵ. Compute the SVD of $\Phi(t_\epsilon)$ and initialize U, V, and the variables ν_i's at t_ϵ.
- For $t \geq t_\epsilon$ solve the differential equations (22) along with the differential equations for U and V,

$$\dot{U} = UH \ , \qquad \dot{V} = VK \ , \tag{24}$$

where H and K are given in (23). If needed, but see below, the singular values can be obtained from the ν_j's by means of equation (20). As already remarked, if we need to approximate only the stability spectra and no information about the growth subspaces is needed, then we can avoid solving the differential equations for V.

3.1.1 Complete algorithm for computing stability spectra

We now give details of how we approximate the stability spectra Σ_{L} and Σ_{ED} and the corresponding leading directions. Since the leading directions, see Facts 2.3, can be recovered from the columns of $\overline{V}$, in Section 3.1.3 we will only clarify how we detect that V has converged in finite precision to its limiting value.

First of all, let us appreciate that we integrate (22), and thus we have always available the quantities

$$\begin{aligned} b_j(t) &= \frac{1}{t}\log \nu_j(t), \qquad j = 1, \ldots, n-1, \\ b_n(t) &= \frac{1}{t}\nu_n(t) \ . \end{aligned} \tag{25}$$

Approximation of Σ_{L}. Now, as far as the approximation of Σ_{L} is concerned, if the system is regular, i.e., the Lyapunov exponents exist as limits, then we will simply approximate the Lyapunov exponents as

$$\lambda_j \approx \lambda_j(T) := \frac{1}{T}\log \sigma_j(T) \ , \ j = 1, \ldots, n \ , \tag{26}$$

for some (large) value of $T > 0$. In terms of the quantities in (25), which is what we really compute, we have

$$\begin{aligned} \lambda_n(T) &= b_n(T) \\ \lambda_j(T) &= \lambda_{j+1}(T) - b_j(T), \quad j = n-1, \ldots, 1 \ . \end{aligned} \tag{27}$$

When the system is not regular, we approximate limsups and liminfs following the same approach of [16]. The idea is to mimic the definition:

$$\lambda_j^s = \limsup_{t\to+\infty} \frac{1}{t}\log \sigma_j(t) = \lim_{\tau\to+\infty} \sup_{t\geq\tau} \frac{1}{t}\log \sigma_j(t).$$

So we pick $\tau > 0$ large, $T >> \tau$, and we use the approximations (as usual, reformulated in terms of the b_j's of (25))

$$\lambda_j^s \approx \max_{\tau\leq t\leq T} \frac{1}{t}\log \sigma_j(t) \ , \qquad \lambda_j^i \approx \min_{\tau\leq t\leq T} \frac{1}{t}\log \sigma_j(t) \ , \tag{28}$$

which will give us approximations to the lim sup's and lim inf's to be used in the approximation of $\Sigma_{\rm L}$.

Approximation of $\Sigma_{\rm ED}$. To approximate $\Sigma_{\rm ED}$, we rely on Steklov averages. First of all, in [14, Theorem 4.6], it is shown that $\Sigma_{\rm ED}$ of (3) is the same as $\Sigma_{\rm ED}$ of the diagonal problem

$$\dot{\Sigma} \;=\; \mathrm{diag}(C)\Sigma\,,\; C = U^T A U\,.$$

At this point, we can proceed as the authors did in [16], that is we use the characterization of $\Sigma_{\rm ED}$ as

$$\Sigma_{\rm ED} = \bigcup_{j=1}^{n} [\alpha_j^S, \beta_j^S], \quad \alpha_j^S = \inf_{t\geq 0} \frac{1}{S}\int_t^{t+S} c_{jj}(s)ds, \quad \beta_j^S = \sup_{t\geq 0} \frac{1}{S}\int_t^{t+S} c_{jj}(s)ds,$$

with S sufficiently large. We finally mimic this characterization to obtain approximate α_j^S and β_j^S by fixing S and $T >> S$, and computing the following quantities

$$\alpha_j^S \approx \min_{0\leq t\leq T} \frac{1}{S}\int_t^{t+S} c_{jj}(s)ds, \qquad \beta_j^S \approx \max_{0\leq t\leq T} \frac{1}{S}\int_t^{t+S} c_{jj}(s)ds. \tag{29}$$

In Section 4, we give some indication of the values of S and T used in practice.

3.1.2 Implementation

Here we give details of how we implemented our continuous SVD method for (3). The modifications needed for the nonlinear problem (1) are straightforward.

We need to integrate (22) along with (24). To form H and K via (23), the simplest thing is to exponentiate $\log(\nu_j)$ to get the ν_j, $j = 1, \ldots, n-1$, and thus the ν_{ij}'s: We have done precisely this.

The ODEs (22) and (24), and (1), are integrated with a 4-th order Runge-Kutta scheme: the 3/8-th rule of [26], henceforth labeled `RK38`. We use variable stepsize by controlling the local error either on the ν_i's or on the orthogonal factor U. Error control is implemented in the standard way explained in [26]. To keep U and V orthogonal, at every step, we replace the approximations returned by `RK38` with the orthogonal factors obtained from the QR factorizations of these approximations, and we use the modified Gram-Schmidt algorithm for this scope.

To update the b_j's at every step, see (25), is simple. Let $\log \nu_j^m$ be the approximation of $\log \nu_j(t_m)$, $j = 1, \ldots, n-1$, obtained with `RK38` at time $t = t_m$:

$$\log \nu_j^m = \log \nu_j^{m-1} + \mathrm{a}pp_j^m, \qquad j = 1, \ldots, n-1,$$
$$\nu_n^m = \nu_n^{m-1} + \mathrm{a}pp_n^m \ ;$$

then, we update b_j from (25).

We recap our continuous SVD algorithm for extracting spectra relative to (3).

Algorithm 3.2 *Continuous SVD Method.*

- Initialize. For t_ϵ given, use `RK38` to integrate $\dot{\Phi} = A(t)\Phi$ up to time t_ϵ. Compute the SVD there: $\Phi(t_\epsilon) = U_0 \Sigma_0 V_0^T$.

- Use U_0 and V_0 as initial conditions for equations (24). The initial conditions for (21) are given by (20) for $t = t_\epsilon$, with $\sigma_j(t_\epsilon)$ being the j-th diagonal element of Σ_0.

- Solve (22) and (24) using (23). While solving these differential equations, update the b_j's.
- Approximate the spectra as explained in Section 3.1.1.

Remark 3.3 For the nonlinear problem (1), we integrate for the trajectory $\phi^t(x_0)$ also using `RK38`. Error control is now performed always on the trajectory, and (additionally) on the variables ν_j's and/or the factor U. Also, we integrate for the trajectory alone for a while before beginning to approximate the spectra; this is very standard practice and it is done to make sure trajectories have passed a possible transient behavior.

3.1.3 Obtaining $\overline{V}$

As we have already remarked, V converges exponentially fast to $\overline{V}$: $\|V(t) - \overline{V}\| \leq ce^{-\alpha t}$, where α is given in (18); notice that for regular systems, α is the distance between the two closest exponents. Obviously, after some finite time $\overline{T}$, $V(t), t \geq \overline{T}$ will be numerically identical to $\overline{V}$. That is, if we let `EPS` be the machine precision, then $\|V(t) - \overline{V}\| \leq$ `EPS`, $\forall t \geq \overline{T}$. In practice, we use the computed approximations of the Lyapunov exponents to predict and refine the estimate of α in this convergence inequality, and from this to predict the value $\overline{T}$ such that $e^{-\alpha \overline{T}} =$ `EPS`. To be precise, we begin estimating α as soon as the distance between two consecutive approximations to V, say V_k and V_{k+1} approximations at t_k and $t_k + h$, is less than or equal to 10 `EPS`: $\alpha = \min_{j=1,\ldots,n-1}(b_j(t_k + h))$, with the b_j's defined in (25). We set $\overline{T} = \frac{1}{\alpha}\log($`EPS`$)$, set $\widetilde{V} = V_{k+1}$ and declare $\widetilde{V}$ to be our approximation to $\overline{V}$. We then continue approximating V up to $\overline{T}$. If the computed approximations to V, until $\overline{T}$, remain within 10 `EPS` of our current $\widetilde{V}$, we accept $\widetilde{V}$ as approximation to $\overline{V}$; otherwise, we update $\widetilde{V}$, $\overline{T}$, and α.

4 Examples

Here we show performance of our SVD algorithms on five problems of different nature. Our goals are to illustrate the following features: (a) ability to approximate both $\Sigma_{\rm L}$ and $\Sigma_{\rm ED}$ for linear and nonlinear problems, (b) possibility, and limitations, of computing directions' growth on linear and nonlinear problems of moderate size, (c) how the algorithms are adapated for computing spectra of Hamiltonian systems.

Example 4.1 This is a regular linear system with LEs: $1, 0, -2, -10$. The coefficient matrix A is such that an orthogonal change of variables Q, reduces it to the sum of a diagonal matrix and a perturbation that goes to zero for t that tends to $+\infty$:

$$Q^T A Q - Q^T \dot{Q} = \mathrm{diag}(1, 0, -2, -10) + B(t), \quad t \geq 0,$$

where

$$B(t) = \begin{pmatrix} \frac{\sin(t)}{t^2+2} + \frac{1}{(t+1)^2} & \frac{\log(t+0.1)}{t+0.2} & \frac{1}{t+0.1} & \frac{t}{t^2+1} \\ \frac{\log(t+0.2)}{t^2+3} & \frac{1}{t+1} + \log\big(\cos\big(\frac{\sqrt{t}}{t+1}\big)\big) & \frac{\sqrt{t}}{t^2+1} & \frac{\cos(t)}{t+1} \\ \frac{1}{t+0.2} & \frac{\log(t+0.1)}{t+2} & \frac{\cos(t)}{t^2+2} + \frac{\sin(t)}{t+1} & \frac{t}{t^2+1} \\ \frac{t}{t^2+1} & \frac{t}{t^2+1} & \frac{\sin(t)}{t+1} & \frac{\log(t+0.1)}{t+1} + \frac{t}{t^2+1} \end{pmatrix}.$$

Since $\lim_{t\to+\infty} \|B(t)\| = 0$, Σ_L is stable ([14], Theorem 3.2, pg.12). The computation of the LEs for this system is easy, hence we chose to omit the results we obtained. Our scope here is to verify how reliable is the computation of the matrix $\overline{V}$. In particular, we want to check the exponential convergence of V to $\overline{V}$ and how well the columns of $\overline{V}$ approximate the growth directions.

Since this is a regular problem, and the minimal distance between the LEs is 1, from (18) the expected value of $\overline{T}$ for which $\|V^T(\overline{T})\overline{V} - I\| \leq$ `EPS` is $\overline{T} \simeq 36$.

We use the continuous SVD Algorithm 3.2, integrating all relevant differential equations with `RK38` and constant stepsize. To monitor the convergence to $\overline{V}$ we proceed as in Section 3.1.3.

In agreement with our expectations, and according to the procedure outlined in Section 3.1.3, after time $\simeq 34$, the computed matrix V does not change

anymore, i.e. we have numerically converged to an orthogonal matrix $\widetilde{V}$. But is this $\widetilde{V}$ a good approximation of $\overline{V}$? To verify this, we check if the exponential behavior of $\Phi\widetilde{v}_i$ is close to λ_i, where $\widetilde{v}_i$ is the i-th column of $\widetilde{V}$. In Figure 1, we plot $\frac{1}{t}\log\|\Phi(t)\widetilde{v}_2\|$ in function of time. The different curves correspond to different stepsizes. Since $\widetilde{v}_2$ is an asymptotic direction, we only expect to be approaching $\lambda_2 = 0$ as time increases. This is clearly achieved, all the more accurately for more accurate computations (smaller stepsize). However, eventually, unavoidable roundoff errors will lead us in the direction of the dominant LE.

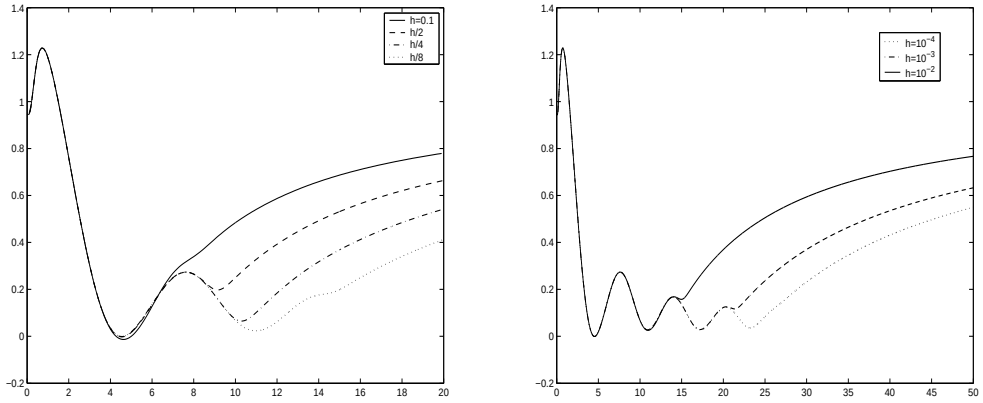

Figure 1. Example 4.1: Plot of $\frac{1}{t}\log\|\Phi(t)\overline{v}_2\|$.

Example 4.2 (Lorenz 40) This non linear system was proposed in [30,31] as a model of an atmospheric variable with values x_j at N equally spaced points around a circle of constant latitude:

$$\dot{x}_k = (x_{k+1} - x_{k-2})x_{k-1} - x_k + 8\ ,\ k = 0, 1, \ldots, N-1\ ,\ (\text{mod } N)\,. \qquad (30)$$

We take the standard value of $N = 40$, for which reported results give 13 positive LEs and Lyapunov dimension δ of the attractor ≈ 27 ([30,31]). [This is the Kaplan-Yorke dimension: $\delta = k + \frac{\sum_{i=1}^{k}\lambda_i}{|\lambda_{k+1}|}$, where k is given by $k:\ \sum_{i=1}^{k}\lambda_i \geq 0$, but $\sum_{i=1}^{k+1}\lambda_i < 0$.] Notice that one LE ($\lambda_{14}$ here) must be equal to zero since the direction of the vector field is preserved. That is, for initial condition x_0 on the attractor, then $\Phi^t(x_0)f(x_0) = f(\varphi^t(x_0))$. For our computations, we obtain x_0 by integrating equation (30) for random IC up to $t = 10^4$. The results displayed in Table 1, 2 and 3 are obtained by means of Algorithm 3.2. To solve the ODEs, we used `RK38` with constant stepsize $h = 10^{-3}$ and $t_\epsilon = h$ as well. Σ_{L} is computed as a continuous spectrum, the values chosen for `T` and τ are shown in the tables as well. For the computation of Σ_{ED} we chose same final time `T`, and time windows of size `S`. In Table 2, we show the obtained values for the LE that we expect to be zero (λ_{14}); in Table 1 and 3 we display the values obtained for the 13th and 15th spectral intervals, respectively. The given

Table 1
Example 4.2: $h = 10^{-3}$

T	S, τ	$\lambda^{\mathrm{i}}_{13}$	$\lambda^{\mathrm{s}}_{13}$	α_{13}	β_{13}
1.0E + 3	1.0E + 2	0.01157	0.04103	−0.02226	0.12155
1.0E + 4	1.0E + 3	0.02427	0.04085	0.00508	0.06914
2.0E + 4	2.0E + 3	0.02427	0.03955	0.01867	0.06556
3.0E + 4	3.0E + 3	0.02713	0.03955	0.02225	0.05721
1.0E + 5	1.0E + 4	0.03379	0.03955	0.02215	0.04176

Table 2
Example 4.2: $h = 10^{-3}$

T	S, τ	$\lambda^{\mathrm{i}}_{14}$	$\lambda^{\mathrm{s}}_{14}$	α_{14}	β_{14}
1.0E + 3	1.0E + 2	−5.04E − 3	5.81E − 3	−4.2E − 2	5.1E − 2
1.0E + 4	1.0E + 3	−4.75E − 3	1.74E − 3	−7.5E − 3	8.1E − 3
2.0E + 4	2.0E + 3	−3.04E − 3	9.54E − 4	−5.3E − 3	5.5E − 3
3.0E + 4	3.0E + 3	−1.04E − 3	6.77E − 4	−3.5E − 3	3.0E − 3
1.0E + 5	1.0E + 4	−5.8E − 4	2.0 − 4	−1.5E − 3	1.5E − 3

Table 3
Example 4.2: $h = 10^{-3}$

T	S,τ	$\lambda^{\mathrm{i}}_{15}$	$\lambda^{\mathrm{s}}_{15}$	α_{15}	β_{15}
1.0E + 3	1.0E + 2	−0.14011	−0.07104	−0.14917	0.02417
1.0E + 4	1.0E + 3	−0.08558	−0.07132	−0.10115	−0.04927
2.0E + 4	2.0E + 3	−0.08558	−0.07426	−0.09440	−0.05603
3.0E + 4	3.0E + 3	−0.08192	−0.07426	−0.09396	−0.06215
1.0E + 5	1.0E + 4	−0.08337	−0.07650	−0.08911	−0.07927

results give numerical evidence of stability of the LEs. Indeed for $\mathtt{T} = 10^4$ and $\mathtt{S} = 10^3$, and for all larger values, the three computed exponential dicothomy intervals are disjoint, and the same is in fact true for all other intervals (all 40 of them). Hence Σ_{ED} is given by the union of n disjoint intervals, each of them containing an interval of Σ_{L}. This and stability of Σ_{ED} imply stability of Σ_{L}. Apparently, as T increases, the Lyapunov intervals tend to points (recall the MET).

Also, we verify convergence of the computed matrix V to a matrix $\widetilde{V}$, using the same technique as in Section 3.1.3. The expected value of time, t, for which convergence to $\widetilde{V}$ has been achieved is such that $e^{t \min_i(\lambda_{i+1}-\lambda_i)} = 10^{-15}$, which is $\simeq 1156$.

Table 4
Example 4.2: $\mathtt{T} = 10^4$, $\|f(x_0)\| = 121.247$

TOL	$\overline{\mathtt{T}}$	$\|\widetilde{V}_u^T f(x_0)\|$	$\|\widetilde{V}_s^T f(x_0)\|$
1.0E − 6	956	2.10E − 1	121.247
1.0E − 9	1416	6.89E − 4	121.247
1.0E − 12	1765	2.64E − 6	121.247

In order to check how reliable is the approximation $\widetilde{V}$ to $\overline{V}$, we did the following experiment. Partition $\overline{V} = (\overline{V}_u, \overline{V}_s)$, where $\overline{V}_u$ comprise the columns of $\overline{V}$ leading to positive LEs, and $\overline{V}_s$ those leading to 0 or negative LEs. Since $\Phi^t(x_0)f(x_0) = f(\varphi^t(x_0))$, and the phase space is compact, then $f(x_0)$ must belong to span$(\overline{V}_s)$. The results displayed in Table 4 are obtained for final $\mathtt{T} = 10^4$. We used Algorithm 3.2 with variable stepsize and error control on the ν_i's. The chosen tolerance for the local error is denoted with TOL in the table. It is clear how the component of $f(x_0)$ in the subspace spanned by the columns of $\widetilde{V}_u$ is going to zero as we increase accuracy in our computations.

Example 4.3 (Rössler) This is the system

$$\dot{x} = -y - z \ , \ \ \dot{y} = x + ay \ , \ \ \dot{z} = b + z(x - c) \ ,$$

and values of interest from [42] are $a = b = 2/10$, $c = 5.7$.

For these values of the parameters, the system has an attractor with a positive LE. As we already noticed in Example 4.2, one LE must be equal to zero. In Tables 5 and 6 we show the intervals of Σ_{L} computed using Algorithm 3.2 using error tolerance $\mathtt{TOL} = \mathtt{1.E-6}$. As indication of the workload, in Table 5, we report also on the number of total steps needed ISTEP, and the CPU time in seconds TIME. From the obtained approximation for Σ_{L}, there is little doubt that Σ_{L} is a point spectrum.

Table 5
Example 4.3: approximation of $[\lambda_1^i, \lambda_1^s]$ with Algorithm 3.2

τ	T	$\lambda_1^{\mathtt{i}}$	$\lambda_1^{\mathtt{s}}$	ISTEP	TIME
1.0E + 3	1.0E + 5	0.07085	0.07633	2.3E + 6	42
1.0E + 4	1.0E + 6	0.07092	0.07374	2.3E + 7	417

Table 6
Example 4.3: approximation of $[\lambda_2^i, \lambda_2^s]$, $[\lambda_3^i, \lambda_3^s]$ with Algorithm 3.2

τ	T	$\lambda_2^{\mathtt{i}}$	$\lambda_2^{\mathtt{s}}$	$\lambda_3^{\mathtt{i}}$	$\lambda_3^{\mathtt{s}}$
1.0E + 3	1.0E + 5	−2.0E − 3	1.8E − 3	−5.40154	−5.38118
1.0E + 4	1.0E + 6	−2.6E − 4	2.2E − 4	−5.39721	−5.39382

Following the procedure described in Section 3.1.3, we observe numerical convergence of V to an orthogonal matrix $\widetilde{V}$. To check how well this matrix $\widetilde{V}$ approximates $\overline{V}$, we proceed as we did in Examples 4.1 and 4.2. In Figure 2 we plot $\frac{1}{t}\log\|\Phi(t)\widetilde{v}_2\|$ for three different $\widetilde{V}$ computed with three different values of TOL. If $\widetilde{V}$ is accurate, we should expect this to be close to $\lambda_2 = 0$. Indeed, the more accurate the computations are (the different value of TOL are indicated in the legend), the longer we remain close to 0.

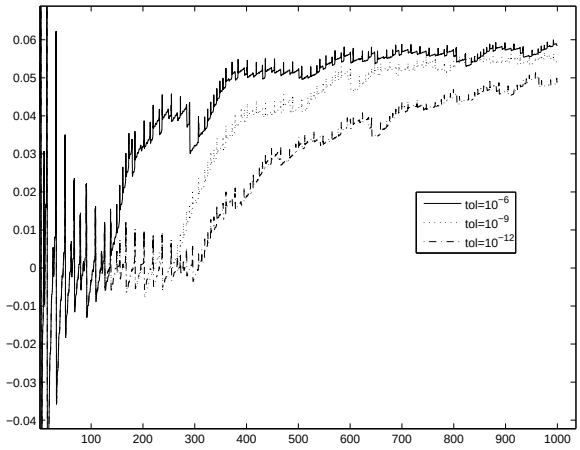

Figure 2. Example 4.3. Plot of $\frac{1}{t}\log\|\Phi\widetilde{v}_2\|$: $\widetilde{V}$ from continuous SVD.

In Table 7 we show the value of the component of $f(x_0)$ in the direction of $\widetilde{v}_1$ and also $\overline{\mathtt{T}}$. The norm of $f(x_0)$ is 8.375.

Table 7
Example 4.3: directions

TOL	$\overline{\mathtt{T}}$	$\lvert\widetilde{v}_1^T f(x_0)\rvert$
1.0E − 6	506	1.5E − 4
1.0E − 7	543	1.8E − 5
1.0E − 8	576	2.1E − 6
1.0E − 9	510	2.2E − 7
1.0E − 10	533	2.6E − 8
1.0E − 11	569	2.7E − 9
1.0E − 12	569	2.7E − 10

Finally, for this example, we include results on the approximation of $\overline{V}$ obtained with the discrete SVD method. We made experiments with constant stepsize $h = 10^{-2}$ and then with $h/2$, using same ICs for the trajectory as we used for the continuous SVD method. Again we observe convergence of V to a matrix $\widetilde{V}$ and an estimate of $\overline{T}$ between 550 and 600. But the main news now is that if we verify how good is this approximation of $\widetilde{V}$ to $\overline{V}$, by verifying if

$\frac{1}{t}\log\|\Phi(t)\widetilde{v}_2\|$ stays close to $\lambda_2 = 0$, we notice that this is indeed true and for a noticeably longer time than in the continuous case; cfr. Figure 3 with Figure 2.

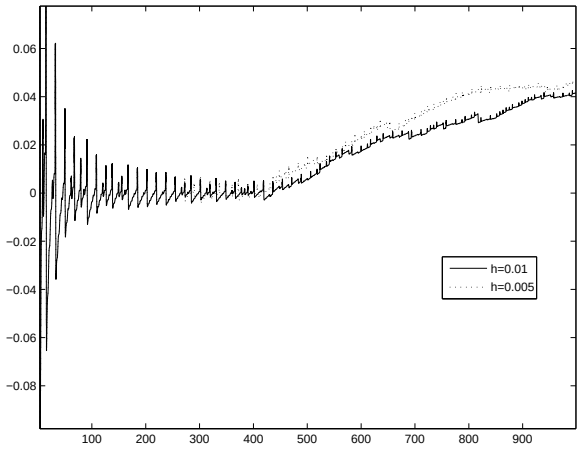

Figure 3. Example 4.3. Plot of $\frac{1}{t}\log\|\Phi\widetilde{v}_2\|$: $\widetilde{V}$ from discrete SVD.

Example 4.4 (Hill equation with quasi periodic forcing) This is the Hill's equation with quasi-periodic potential

$$\ddot{x} + (a^2 + bp(t))x = 0, \tag{31}$$

where $p = p(t) = \cos(t) + \cos(\gamma t)$, $\gamma = \frac{1+\sqrt{5}}{2}$, and a and b are parameters. We write this as the Hamiltonian system

$$\dot{x} = y \ , \ \ \dot{y} = (a^2 + bp(t))x \ .$$

Hamiltonian systems have the property that Σ_{L} and Σ_{ED} are symmetric with respect to the origin. More specifically, the singular values of the linearized problem (which is symplectic), arise as pairs σ and $1/\sigma$. Hence we can just compute half of the spectrum to recover all of it.

The Hill equation with quasi periodic potential is studied in [9]. There, the authors rewrite (31) as a 4 dimensional non linear dynamical system, by introducing the new variables $t, \gamma t$. Their purpose is to study the linear behavior of the invariant 2-torus $\mathbb{T}^2 \times \{0, 0\}$. However, we are just interested in the computation of the Lyapunov spectrum of (31) and to this end it suffices to study the linear system. In Table 8 we show approximation to λ_1 for $a = 0.413$ and $b = 0.4$. We use Algorithm 3.2 together with `RK38` with variable stepsize. We check the error on the exponents (`IPAR` $= 0$) or on the matrix U (`IPAR` $= 1$) and we choose the tolerance for the local error equal to `TOL` $= 10^{-12}$. By looking at the intervals in the Table, the Lyapunov spectrum appears to be a point spectrum.

Table 8
Example 4.4: approximation of $[\lambda_1^i, \lambda_1^s]$ for $a = 0.413$, $b = 0.4$.

τ	T	IPAR	$\lambda_1^{\mathtt{i}}$	$\lambda_1^{\mathtt{s}}$	ISTEP	IREJ	TIME
1.0E + 3	1.0E + 4	0	0.00971	0.01180	1.2E + 6	1.1E + 4	7.3
1.0E + 3	1.0E + 4	1	0.00971	0.01180	3.0E + 6	2.0E + 2	8.9
1.0E + 4	1.0E + 5	0	0.01017	0.01040	1.2E + 6	1.1E + 4	42.3
1.0E + 4	1.0E + 5	1	0.01018	0.01040	3.0E + 7	1.9E + 3	52
1.0E + 5	1.0E + 6	0	0.01023	0.01025	3.9E + 7	1.1E + 6	301
1.0E + 5	1.0E + 6	1	0.01023	0.01025	3.0E + 8	1.9E + 4	299

Figure 4 shows plots of λ_1 as a function of a^2 for $b = 0.2$ and $b = 0.6$ (on the left and right, respectively). These plots are in agreement with those in [9].

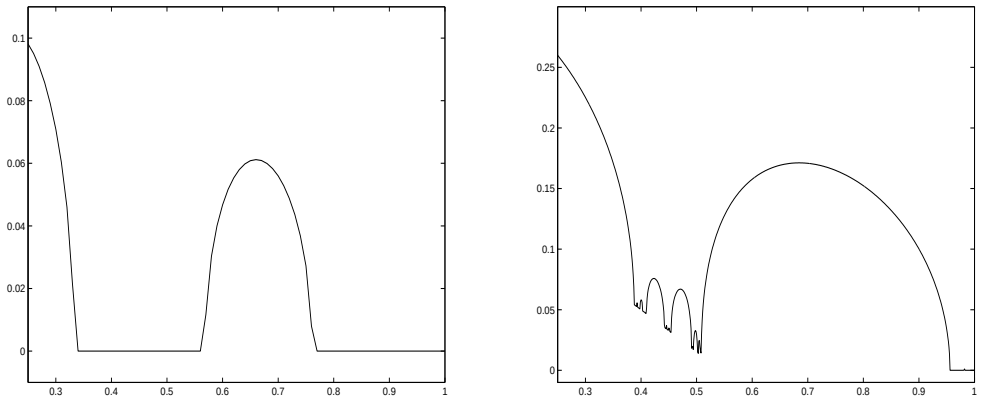

Figure 4. Example 4.4: Plot of the dominant LE as a function of a^2.

Example 4.5 (Hènon-Heiles) This Hamiltonian system is often used as test problem for computation of Lyapunov exponents; e.g., see [29]. We simply compute the reduced SVD corresponding to half of the spectrum (this also avoids potential difficulties with the double Lyapunov exponent at 0).

We have the Hamiltonian

$$H(q_1, q_2, p_1, p_2) = \frac{1}{2}(p_1^2 + p_2^2) + \frac{1}{2}(q_1^2 + q_2^2) + q_1^2 q_2 - \frac{1}{3} q_2^3 \,,$$

and the associated dynamical system $\dot{x} = J\nabla H(x(t))$. That is, we have

$$\begin{aligned} \dot{q}_1 &= p_1 \\ \dot{q}_2 &= p_2 \\ \dot{p}_1 &= -q_1(1 + 2q_2) \\ \dot{p}_2 &= -q_2 - q_1^2 + q_2^2. \end{aligned} \tag{32}$$

We must make sure that the integration for the solution of (32) be done with a symplectic scheme, otherwise the aforementioned symmetries in the spectra will not be maintained. On the other hand, extraction of the spectral information can be done as explained in Section 3.

So, for (32), and more generally for Hamiltonian problems, we have implemented the following strategy. Suppose h is the stepsize with which we want to proceed for the SVD methods, to go from t_i to t_{i+1}. Then, we integrate (32) with the 4th order Gauss-Runge-Kutta scheme, which is symplectic, with stepsize $h/2$, obtaining (4th order accurate) approximations to the solution at $t_i, t_{i+1/2}, t_{i+1}$. [The scheme is implicit, and we solve the nonlinear systems by an inexact Newton's method keeping the Jacobian evaluated –and factored– at the initial value]. We then use these approximations to setup the linearized problem, and use the "classical" 4th order Runge-Kutta scheme (which requires information at the initial point, mid-point, and final point), `RK4`, to solve the differential equations needed for the SVD method. The end result is a 4th order method.

For (32), for comparison with the results in [29], we performed two experiments with the initial conditions (IC) in Table 9 with q_1 always given by 0:

Table 9
Initial conditions.

IC	q_2	p_2	p_1
1	0.20	0.14	0.442417
2	0.25	0.30	0.328506

In Table 10 we show the approximate Lyapunov exponents obtained for final time `T`. To solve the ODEs, we implemented `RK4` with constant stepsize $h = 10^{-2}$. Our results compare well with the ones in [29].

5 Conclusions

To approximate Lyapunov and Dichotomy spectra, Σ_{L} and Σ_{ED}, we have presented algorithms based on the SVD of the linearized system. These methods are backed up by a solid theory for problems with stable and distinct Lyapunov exponents (integrally separated problems). We have shown how to overcome possible pitfalls of SVD methods, and shown practical performance of the techniques on several examples.

In our experience, SVD based methods handle satisfactorily small to moderate sized problems (say, of dimension up to 100), for approximation of Σ_{L} and Σ_{ED}.

Table 10

IC	T	λ_1	λ_2
1	1.0E2	1.3E − 1	4.9E − 2
1	1.0E3	7.6E − 2	7.3E − 3
1	1.0E4	4.5E − 2	1.0E − 3
1	1.0E5	3.9E − 2	1.1E − 4
2	1.0E2	6.8E − 2	6.2E − 2
2	1.0E3	5.8E − 2	7.2E − 3
2	1.0E4	6.0E − 2	9.0E − 4
2	1.0E5	4.9E − 2	1.1E − 4

For large problems, more effort is needed to obtain efficient techniques. One of the main appeals of SVD techniques is the ability to recover not just the spectra, but also the geometric information on the growth subspaces. As far as we know, our methods are the first ones who have been able to approximate these subspaces.

Careful comparison with QR methods has not been carried out, but it is certainly of interest and may be the subject of future investigation.

6 Appendix

Here below are the tableaux of the Runge-Kutta schemes we used. The last two rows in the table for RK38 are the additional stage and weights needed to form the lower order scheme used for error control.

0	0	0	0	0	0
$\frac{1}{3}$	$\frac{1}{3}$	0	0	0	0
$\frac{2}{3}$	$-\frac{1}{3}$	1	0	0	0
1	1	−1	1	0	0
	$\frac{1}{8}$	$\frac{3}{8}$	$\frac{3}{8}$	$\frac{1}{8}$	0
1	$\frac{1}{8}$	$\frac{3}{8}$	$\frac{3}{8}$	$\frac{1}{8}$	0
	$\frac{1}{12}$	$\frac{1}{2}$	$\frac{1}{4}$	0	$\frac{1}{6}$

3/8-th Runge-Kutta 4-3 pair: RK38.

0	0	0	0	0
$\frac{1}{2}$	$\frac{1}{2}$	0	0	0
$\frac{1}{2}$	0	$\frac{1}{2}$	0	0
1	0	0	1	0
	$\frac{1}{6}$	$\frac{2}{6}$	$\frac{2}{6}$	$\frac{1}{6}$

$1/2 - \sqrt{3}/6$	$1/4$	$1/4 - \sqrt{3}/6$
$1/2 + \sqrt{3}/6$	$1/4 + \sqrt{3}/6$	$1/4$
	$1/2$	$1/2$

Classical 4th order Runge-Kutta and Gauss Runge-Kutta of order 4.

References

[1] H. D. I. Abarbanel, R. Brown, and M. B. Kennel. Local Lyapunov exponents computed from observed data. *Journal of Nonlinear Science*, 2:343–366, 1992.

[2] L. Ya. Adrianova. *Introduction to Linear Systems of Differential Equations*. Translations of Mathematical Monographs 146. AMS, Providence R.I., 1995.

[3] L. Arnold. *Random Dynamical Systems*, 2nd Edition. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003.

[4] L. Arnold and V. Wihstutz Eds. *Lyapunov Exponents. Proceedings, Bremen 1984*. Springer-Verlag, Berlin, 1986. Lecture Notes in Mathematics 1186.

[5] L. Arnold, H. Crauel, and J.P. Eckmann Eds. *Lyapunov Exponents. Proceedings, Oberwolfach 1990*. Springer-Verlag, Berlin, 1991. Lecture Notes in Mathematics 1486.

[6] G. Benettin, G. Galgani, L. Giorgilli, and J. M. Strelcyn. "Lyapunov exponents for smooth dynamical systems and for Hamiltonian systems; a method for computing all of them. Part I: Theory", and "...Part II: Numerical applications", *Meccanica*, 15:9–20,21–30, 1980.

[7] A. Bojanczyk, G. Golub, and P. Van Dooren. "The periodic Schur form. Algorithms and applications". In *Proceedings SPIE Conference*, pages 31–42, San Diego, 1992.

[8] A. Bojanczyk, M. Ewerbring, F. Luk and P. Van Dooren. "An accurate product SVD algorithm". In *SVD and signal processing, II. Algorithms, Analysis and Applications*, Ed. R.J. Vaccaro, Eslevier Publ., pages 113–131, 1991.

[9] H. Broer, C. Simó. "Hill's equation with quasi–periodic forcing: resonance tongues, instability pockets and global phenomena". In *Bul. Soc. Bras. Mat.*, 29: 253–293, 1998.

[10] F. Christiansen and H. H. Rugh. Computing Lyapunov spectra with continuous Gram-Schmidt orthonormalization. *Nonlinearity*, 10:1063–1072, 1997.

[11] P. Constantin and C. Foias. Global Lyapunov exponents, Kaplan-Yorke formulas and the dimension of the attractors for 2D Navier-Stokes equations. *Comm. Pure Appl. Math.*, 38:1–27, 1985.

[12] W. A. Coppel. *Dichotomies in Stability Theory*. Springer-Verlag Lecture Notes in Mathematics Vol. 629, 1978.

[13] Dieci L., Eirola T.. "On smooth decompositions of matrices", *SIAM J. Matrix Anal. Appl.*, Vol. 20, No. 3 (1999), pp. 800–819.

[14] L. Dieci, and C. Elia. "The singular value decomposition to approximate spectra of dynamical systems. Theoretical aspects", to appear in *Journal of Differential Equations*, 2006.

[15] L. Dieci, C. Elia C., and L. Lopez. "Smooth SVD on the Lorentz group with application to computation of Lyapunov exponents", *J. Comput. & Applied Math.*, Vol 164-165 (2004), pp. 255-264.

[16] L. Dieci and E. Van Vleck. Lyapunov spectral intervals: theory and computation. *SIAM J. Numer. Anal.*, 40:516–542, 2003.

[17] L. Dieci, R. D. Russell, and E. S. Van Vleck. "On the computation of Lyapunov exponents for continuous dynamical systems", *SIAM J. Numer. Anal.*, 34:402–423, 1997.

[18] L. Dieci and E.S. Van Vleck, "Lyapunov and Sacker-Sell spectral intervals," to appear in *J. of Dynamics and Differential Equations*, 2006.

[19] L. Dieci and E.S. Van Vleck, "`LESLIS` and `LESLIL`: Codes for approximating Lyapunov exponents of linear systems", Technical Report (2004): www.math.gatech.edu/∼dieci.

[20] Roberta Fabbri, Russell Johnson, and Francesca Mantellini. A nonautonomous saddle-node bifurcation pattern. *Stoch. Dyn.*, 4(3):335–350, 2004.

[21] K. Geist, U. Parlitz, and W. Lauterborn. "Comparison of different methods for computing Lyapunov exponents", *Prog. Theor. Phys.*, 83:875–893, 1990.

[22] I. Goldhirsch, P. L. Sulem, and S. A. Orszag. Stability and Lyapunov stability of dynamical systems: a differential approach and a numerical method. *Physica D*, 27:311–337, 1987.

[23] J.M. Greene and J-S. Kim. "The calculation of Lyapunov spectra", *Physica D*, 24:213–225, 1987.

[24] W. Hahn. *Stability of Motion*. Springer-Verlag, Berlin-Heidelberg, 1967.

[25] M.T. Heath, A.J. Laub, C.C. Paige and R.C. Ward. "Computing the Singular Vlaue Decomposition of a product of two matrices", *SIAM J. Sci. Stat. Comput.*, 7:1147–1159, 1986.

[26] E. Hairer, S. P. Nœrsett, and G. Wanner. *Solving Ordinary Differential Equations I*. Springer-Verlag, Berlin-Heidelberg, 1993. Second edition.

[27] Russell Johnson and Francesca Mantellini. A nonautonomous transcritical bifurcation problem with an application to quasi-periodic bubbles. *Discrete Contin. Dyn. Syst.*, 9(1):209–224, 2003.

[28] R. A. Johnson, K. J. Palmer, and G. Sell. "Ergodic properties of linear dynamical systems", *SIAM J. Mathem. Analysis*, 18:1–33, 1987.

[29] B. Leimkuhler and E. Van Vleck. Orthosymplectic integration of linear Hamiltonian systems. *Numer. Math.*, Vol. 77 (1997), pp. 269–282.

[30] E. Lorenz. "Predictability. A problem partly solved", in Proceedings on predictability held at ECMWF on 4-8 September 1995, pp. 1-18. [Available from ECMWF, Shinfield Park, Reading, Berkshire RG29AX, e.kooij@ecmwf.int] 1–18, 1996.

[31] E. Lorenz and K. Emmanuel. "Optimal sites for supplementary weather observations: Simulations with a small model", *J. Atmos. Sci.*, 55: 399–414, 1998.

[32] K. Lust. "Improved numerical Floquet multipliers", *Internat. J. Bifur. Chaos Appl. Sci. Engrg.*, 11(9):2389–2410, 2001.

[33] A. Lyapunov. "Problém géneral de la stabilité du mouvement", *Int. J. Control*, 53:531–773, 1992.

[34] S. Oliveira and D.E. Stewart. Exponential splitting of products of matrices and accurately computing singular values of long products. *LAA*, 309:175–190, 2000.

[35] Oseledec V.I., "A multiplicative ergodic theorem. Lyapunov characteristic numbers for dynamical systems", Trans. Moscow Mathem. Society, Vol. 19 (1968), pp. 197–231.

[36] E. Ott. *Chaos in Dynamical Systems, 2nd Edition.* University Press, Cambridge, 2002.

[37] E. Ott, T. Sauer, and J. Yorke. *Coping with Chaos.* Wiley, New York, 1994.

[38] K.J. Palmer, Shadowing in Dynamical Systems, Mathematics and Its Applications, Vol. 501, Kluwer, 2000.

[39] T.S. Parker and L.O. Chua. *Practical Numerical Algorithms for Chaotic Systems.* Springer, New York, 1989.

[40] D. Ruelle. Characteristic exponents and invariant manifolds in hilbert space. *Annals of Mathematics*, 115:243–290, 1982.

[41] R. J. Sacker and G. R. Sell, "A spectral theory for linear differential systems," *J. Diff. Eqn.* **7** (1978), pp. 320–358.

[42] M. Sano and Y. Sawada. Measurement of the Lyapunov spectrum from a chaotic time series. *Physical Review Letter*, 55, 1082, 1985.

[43] D.E. Stewart, "A new algorithm for the SVD of a long product of matrices and the stability of products", *ETNA* **5** (1997), pp. 29–47.

[44] G.W. Stewart, "On graded QR decompositions of products of matrices", *ETNA* **3** (1995), pp. 39–49.

[45] W. Tucker. *The Lorenz attractor exists*. PhD thesis, Uppsala University, Department of Mathematics, 2003.

[46] S. Wiggins, *Normally Hyperbolic Invariant Manifolds in Dynamical Systems*, Springer-Verlag, New York (1994).

[47] K. Wright. Differential equations for the analytical singular value decomposition of a matrix. *Numer. Math.*, 63:283–295, 1992.