Sensitivity Analysis of Chaotic Problems using a Fourier Approximation of the Least-Squares Adjoint

Anthony Ashley* and Jason E. Hicken†
Rensselaer Polytechnic Institute, Troy, New York, 12180

In general, chaotic dynamical systems produce ill-conditioned direct and adjoint sensitivity equations, so these conventional sensitivity analysis methods do not provide useful derivatives for design optimization of chaotic systems. The least-squares shadowing (LSS) method is a recently proposed regularization of the sensitivity equations that leads to useful derivative information; however, the LSS method requires the solution of a second-order boundary-value problem in time that is prohibitively expensive for large-scale simulations. The primary contribution of this work is the observation that the LSS primal and dual solution spaces do not need to be the same, and they can be different dimensions when discretized. For example, we investigate the approximation of the LSS adjoint using a Fourier sine series, and we obtain encouraging results on a model problem based on the Lorenz ordinary differential equation. For the Lorenz problem, the dimensions of the LSS system are reduced by an order of magnitude.

I. Introduction

Chaotic dynamical systems are ubiquitous in science and engineering. Examples from aerospace engineering include the turbulent combustion in an engine or the aeroelastic response of a helicopter blade. In the past, simulations of such systems may have eliminated the intrinsic chaotic behavior by introducing simplifications or closure models; however, as computational power increases, there is a growing interest in using high-fidelity models that more accurately capture the dynamics. Here we have in mind the trend towards unsteady Reynolds-averaged Navier-Stokes or large-eddy simulations of turbulence.

This sets the stage for the following question: can we optimize engineering systems that are governed by chaotic dynamics? Given the success of partial-differential equation (PDE)-governed optimization in structural,[1] aerodynamic,[2–8] and aerostructural design,[9–12] it is natural to assume that we can apply similar methods to chaotic systems, albeit at an increased computational cost. Unfortunately, the blind application of PDE-governed optimization to chaotic systems fails.

In the context of optimization, the difficulty presented by chaotic systems is their sensitive dependence on parameters. Consider the time-averaged sectional lift coefficient on a thin airfoil at high Reynolds number. If we run an unsteady flow simulation on the airfoil, over a range of angles of attack ($\Delta\alpha \ll 1$), we will find that the lift slope is approximately 2π, as expected from thin-airfoil theory; however, the lift data will be scattered about the line, much like experimental data. The scattering is due to the sensitive dependence between the flow solution and the angle of attack. Two simulations, run with identical parameters except for a small difference in their angles of attack, will eventually produce noticeably different flow fields. Consequently, their time-averaged lifts will differ, and the lift data will appear noisy.

This "noise" leads to the failure of conventional sensitivity analysis methods to produce useful derivatives of time-averaged quantities. In the case of the lift-slope example, these methods compute derivatives that are orders of magnitude too large, because they see the highly oscillatory noise rather than the overall slope of 2π. In short, these methods "cannot see the forest for the trees."

Sensitivity analysis methods have been proposed to estimate the derivatives of time-averaged outputs that depend on chaotic states, including the ensemble[13–15] and least-squares shadowing[16] (LSS) sensitivity methods. The LSS adjoint, in particular, has shown great promise; however, its computational cost is high,

*Graduate Student, Department of Mechanical, Aerospace, and Nuclear Engineering, Student Member AIAA
†Assistant Professor, Department of Mechanical, Aerospace, and Nuclear Engineering, Member AIAA

as it requires the solution of a second-order boundary-value problem in time. In Liao,[17] this cost was partially addressed by developing a block-based lower-upper factorization of the LSS system. In the present work, we present a method in which we allow the primal and dual solutions to be of different dimensions. Specifically, we represent the adjoint solution for our model problem using a Fourier sine series, although this choice is not unique.

The paper is organized as follows. We begin by introducing the model problem for sensitivity analysis and reviewing the conventional adjoint in Section II. The LSS adjoint and our proposed Fourier approximation to it are presented in Section III. Numerical results are presented in Section IV, and conclusions arising from the results are discussed in Section V.

II. Model problem and review of the conventional adjoint

A. The generic and model problems

Our intent is to find the sensitivity of some objective function, $\mathcal{J}$, to a design variable. The objective function of interest will be governed by a system of equations that exhibits chaotic behavior. We represent the chaotic dynamical system with the generic initial-value problem (IVP)

$$\begin{aligned} \mathcal{R}(u,\nu,t) \equiv \frac{du}{dt} - \mathcal{F}(u,\nu,t) = 0, \qquad \forall t \in [0,T], \\ u(0) = u_{\mathsf{IC}}, \end{aligned} \tag{1}$$

where $\mathcal{R}$ is the nonlinear residual, and the symbol $\nu \in \mathbb{R}^n$ appearing in $\mathcal{F}$ denotes a control/design variable. Thus, one can consider the solution to (1) to be a function of both time and the parameter ν, i.e. $u = u(t,\nu)$.

Suppose we are interested in minimizing some functional that depends on the solution $u(t,\nu)$. Specifically, consider the time-integrated quantity

$$\mathcal{J}(u,\nu) = \int_0^T \mathcal{G}(u,\nu,t)\, dt, \tag{2}$$

where $\mathcal{G}$ is a nonlinear function. In the fluid mechanics applications we are interested in, $\mathcal{J}(u,\nu)$ may be a time-averaged force or moment.

In order to both demonstrate chaotic behavior and validate our method, we select the Lorenz system for investigation, which is defined by

$$\mathcal{F} = \begin{pmatrix} \sigma(y-x) \\ x(\rho - z) - y \\ xy - \beta z \end{pmatrix}; \tag{3}$$

the state variables of (3) are $u = [x,y,z]^T$, and the parameters of (3) are $\nu = [\sigma,\rho,\beta]^T$. The system has initial conditions $\boldsymbol{x}(0) = \boldsymbol{x}_0$.

The objective function selected for the Lorenz model problem is

$$\mathcal{J} = \frac{1}{T}\int_0^T z \,\mathrm{dt}, \tag{4}$$

as it is well studied and behaves in a consistent manner — the gradient $\partial\mathcal{J}/\partial\rho$ remains relatively constant at a value of 0.96.[13] Note that (4) is in the form (2) with $\mathcal{G}(u,\nu,t) = z/T$.

The sensitive dependence of chaotic systems to initial conditions is fairly well known. What is less widely known is that these systems can be sensitive to changes in other parameters. Figure 1 illustrates this sensitivity. In both runs displayed, the initial conditions and the parameters β and σ were fixed, while the ρ parameters differ by 10^{-10}. After only 15 units of time, the values of the x state variable are drastically different.

B. The conventional adjoint

When there are many inputs and few outputs, the *de facto* sensitivity analysis method in PDE-governed gradient-based optimization is the standard adjoint method. We are interested in such a scenario — although

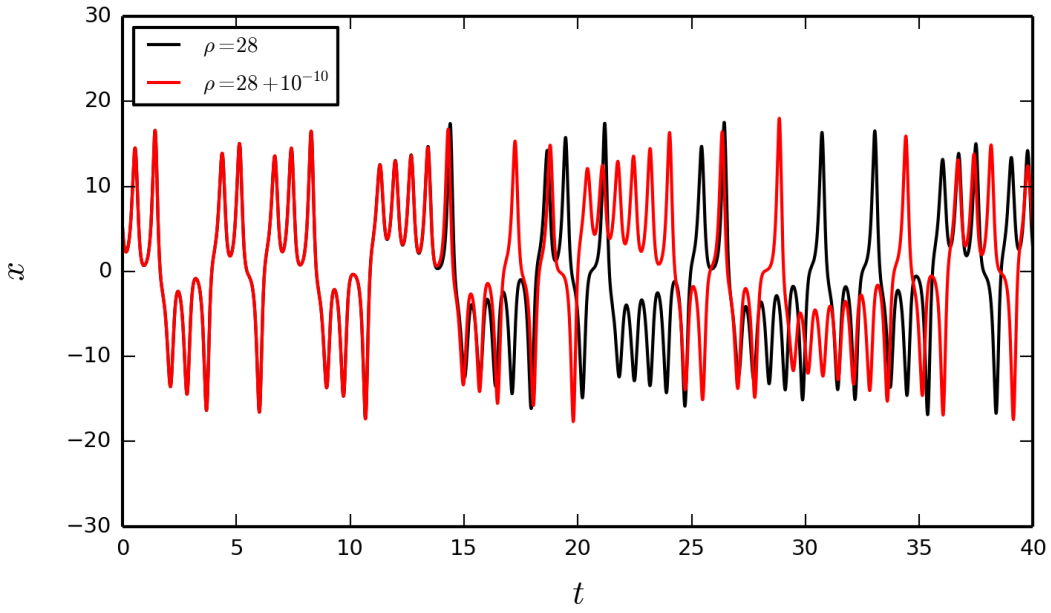

Figure 1. A demonstration of the characteristics of the Lorenz system. Two solutions of the Lorenz system are shown, with the x state variable plotted against time t. Despite being solved with almost identical values of the design variable ρ, the trajectories rapidly diverge. β and σ were set to their typical values of $8/3$ and 10, respectively.

our model problem only has three inputs (ρ, σ, and β) — as we seek a method that will work in aerospace applications that have hundreds or even thousands of design variables. We now review the conventional adjoint for completeness and demonstrate it on the Lorenz problem.

In order to minimize $\mathcal{J}(u, \nu)$ with respect to ν using efficient gradient-based methods, we need the gradient $\nabla_\nu \mathcal{J}$, which is a total derivative. In the case of the functional defined above, $\mathcal{J}$ may depend directly on ν, via its appearance in $\mathcal{G}(\cdot, \nu, \cdot)$, or indirectly through $u(t, \nu)$. In order to account for the indirect dependence, we introduce the Lagrangian

$$\mathcal{L}(u, \psi, \nu) = \mathcal{J}(u, \nu) + \int_0^T \psi^T \mathcal{R}(u, \nu, t)\, dt + \left[\psi^T (u - u_{\mathsf{IC}})\right]_{t=0}, \tag{5}$$

where ψ denotes the adjoint. As long as $\mathcal{R}(u, \nu, t) = 0$ and $u(0) = u_{\mathsf{IC}}$, we have that $\mathcal{L} = \mathcal{J}$; however, their partial derivatives are not necessarily equal.

The gradient we want is given by

$$\nabla_\nu \mathcal{J} \equiv \frac{\partial \mathcal{L}}{\partial \nu} = \int_0^T \frac{\partial \mathcal{G}}{\partial \nu}\, dt + \int_0^T \psi^T \frac{\partial \mathcal{R}}{\partial \nu}\, dt. \tag{6}$$

However, we cannot compute this without first solving for ψ. To find an equation for ψ, we take the first variation of $\mathcal{L}$ with respect to u and set the result to zero:

$$\begin{aligned}
\mathcal{L}(\delta u, \psi, \nu) &= \int_0^T \mathcal{G}(\delta u, \nu)\, dt + \int_0^T \psi^T \mathcal{R}(\delta u, t; \nu)\, dt + \left[\psi^T \delta u\right]_{t=0} \\
&= \int_0^T \frac{\partial \mathcal{G}}{\partial u} \delta u\, dt + \int_0^T \psi^T \frac{\partial \mathcal{R}}{\partial u} \delta u\, dt + \left[\psi^T \delta u\right]_{t=0} \\
&= \int_0^T \left[\frac{\partial \mathcal{G}}{\partial u}^T - \frac{d\psi}{dt} - \frac{\partial \mathcal{F}}{\partial u}^T \psi\right]^T \delta u\, dt + \left[\psi^T \delta u\right]_{t=T} \\
&= 0,
\end{aligned}$$

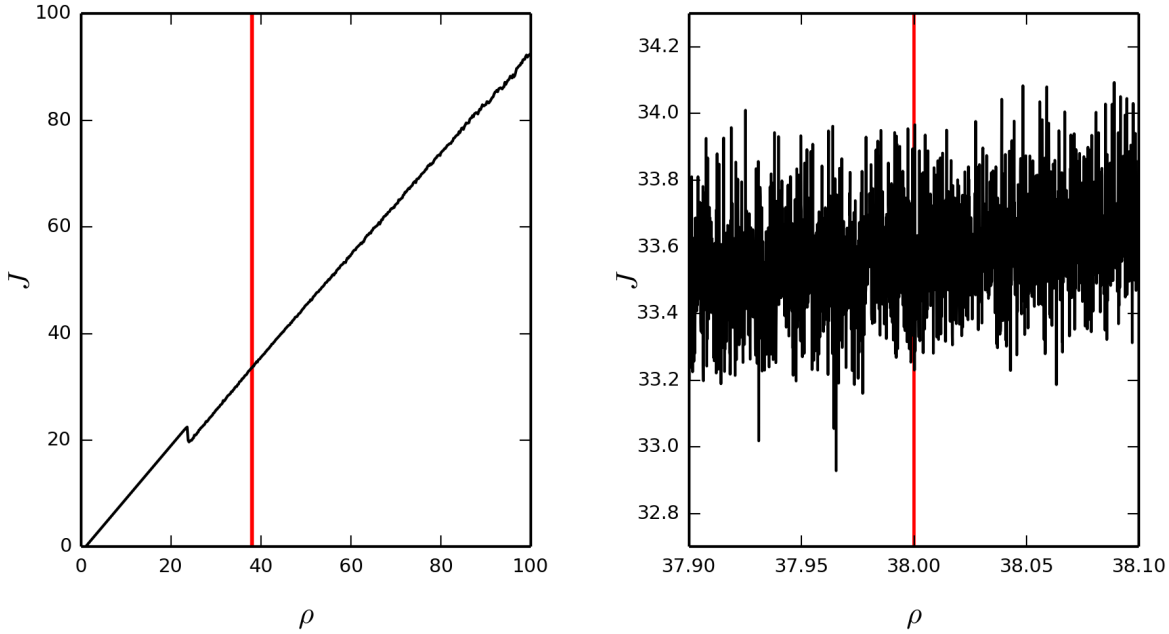

Figure 2. A plot of the objective function $\mathcal{J} = 1/T \int_0^T z$ versus ρ. The left plot shows the $\mathcal{J}$ over a range of $\rho \in [0, 100]$, while the right plot shows J over a range of $\rho \in [37.9, 38.1]$. Overlaid in red is a representative example of the gradient calculated with the conventional adjoint. Immediately visible is the problematic nature of these gradients.

where (1) and integration by parts (in time) are used to arrive at the second last line. Since $\mathcal{L}(\delta u, \psi, \nu) = 0$ must be true for all variations δu, we conclude that

$$\begin{aligned} -\frac{d\psi}{dt} - \frac{\partial \mathcal{F}}{\partial u}^T \psi + \frac{\partial G}{\partial u}^T &= 0, \qquad \forall t \in [0, T], \\ \psi(T) &= 0. \end{aligned} \tag{7}$$

The equation (7) is the conventional adjoint differential equation for the objective (2) and IVP (1).

In a chaotic problem, however, the conventional adjoint grows unbounded with time, making the derivative information unusable. This unbounded growth is related to the sensitive dependence of the state on the initial condition. Or, restated another way, the time-integrated objective function $\mathcal{J}$ looks noisy for finite T, and the adjoint-based gradient accurately reflects this noise, since the derivative (and the adjoint) of a noisy function is large.

Returning to our model problem of the Lorenz system (3) and our selected objective function (4), we can apply the conventional adjoint approach in an attempt to obtain gradient information. Plotted in Figure 2 are discrete solutions of the objective function for a range of values of ρ, along with its gradient with respect to ρ, calculated using the conventional adjoint method. Clearly the adjoint-based derivative does not capture the large-scale trend in the objective function and instead reflects the noise.

III. Least-squares shadowing and its Fourier approximation

A. Least-squares shadowing adjoint

To overcome the instability of the conventional adjoint, the least-squares shadowing (LSS) adjoint method modifies the Lagrangian (5) in a couple ways. In order to describe the LSS adjoint, it is helpful to show how the conventional adjoint is related to a fully-constrained optimization problem and associated saddle-point

problem. Taking (5) and linearizing about $\bar{u}$ using perturbations u', we obtain

$$\mathcal{L}(\bar{u}, u', \psi, \nu) = \mathcal{J}'(u') + \int_0^T \psi^T \mathcal{R}'(u', t)\, dt + \psi^T u', \tag{8}$$

where $\mathcal{R}'$ is given by

$$\mathcal{R}'(u', t) \equiv \left[\frac{d}{dt} - \left. \frac{\partial \mathcal{F}}{\partial u} \right|_{\bar{u}} \right] u' = 0, \tag{9}$$

and $\mathcal{J}'$ is defined as

$$\mathcal{J}'(u') = \int_0^T \mathcal{G}(\bar{u})\, dt + \int_0^T \frac{\partial \mathcal{G}}{\partial u} u'\, dt.$$

If we take the derivative of $\mathcal{L}$ in (8) with respect to u', we obtain the standard adjoint.

In the above definition, $\bar{u}$ is the frozen, or reference, state: it is the state about which we want to evaluate the gradient. The perturbation to the state is the function $u'(t)$. To simplify notation, the dependence of various quantities on ν is not shown in the following derivation.

The issue with the conventional adjoint is that the differential equation (DE) associated with $\mathcal{R}'(u', t)$ is linearly unstable. To regularize this differential equation, the LSS formulation requires the perturbation, u', to satisfy the linearized differential equation without the initial condition. However, with the initial condition removed, the linearized differential equation is also ill-posed; this necessitates the addition of regularization.

The regularization ensures that, of all possible perturbations to the state, u' is the one that i) satisfies the linearized differential equation and ii) minimizes the composite objective consisting of the linearized objective and the L^2 norm of u'. Thus, u' is the solution to the DE-constrained optimization problem

$$\begin{aligned} &\min_{u'} \ \mathcal{J}'(u') + \frac{1}{2}\|u'\|^2 \\ &\text{s.t.} \qquad \mathcal{R}'(u', t) \equiv \left[\frac{d}{dt} - \left. \frac{\partial \mathcal{F}}{\partial u} \right|_{\bar{u}} \right] u' = 0, \qquad \forall t \in [0, T], \end{aligned} \tag{10}$$

where $\|u'\|^2 = \int_0^T (u')^2\, dt$. Based on the optimization problem (10), we can construct the associated Lagrangian:

$$\mathcal{L}'(u', \psi) = \mathcal{J}'(u') + \frac{1}{2}\|u'\|^2 + \int_0^T \psi^T \mathcal{R}'(u', t)\, dt. \tag{11}$$

The LSS adjoint equation can be found by taking the first-variation of $\mathcal{L}'$ with respect to u', which produces the LSS problem

$$\begin{aligned} -\frac{d\psi}{dt} - \frac{\partial \mathcal{F}}{\partial u}^T \psi + \frac{\partial G}{\partial u}^T + u' = 0, \qquad \forall t \in [0, T], \\ \psi(0) = 0, \qquad \text{and} \qquad \psi(T) = 0. \end{aligned} \tag{12}$$

This DE is almost identical to the ill-posed one we found earlier, i.e. (7), except for the presence of u'. We can eliminate u' from the adjoint DE by recalling that $\mathcal{R}'(u', t) = 0$, that is, the perturbation is in the null space of the linearized forward problem. Thus, applying the linear operator corresponding to $\mathcal{R}'$ to the above equation, we get the differential equation

$$\begin{aligned} \left[\frac{d}{dt} - \frac{\partial \mathcal{F}}{\partial u} \right] \left[-\frac{d\psi}{dt} - \frac{\partial \mathcal{F}}{\partial u}^T \psi + \frac{\partial G}{\partial u}^T \right] = 0, \qquad \forall t \in [0, T], \\ \psi(0) = 0, \qquad \text{and} \qquad \psi(T) = 0. \end{aligned} \tag{13}$$

Equation (13) is a linear boundary-value problem (BVP) in time; therefore, two boundary conditions (in time) are necessary, and the conditions at $t = 0$ and $t = T$ are consistent with this requirement.

For completeness, we expand the BVP fully to obtain

$$-\frac{d^2\psi}{dt^2} - \frac{d}{dt}\left[\frac{\partial \mathcal{F}}{\partial u}^T \psi \right] + \frac{\partial \mathcal{F}}{\partial u}\frac{d\psi}{dt} + \frac{\partial \mathcal{F}}{\partial u}\frac{\partial \mathcal{F}}{\partial u}^T \psi = -\frac{\partial^2 \mathcal{G}}{\partial t \partial u} + \frac{\partial \mathcal{F}}{\partial u}\frac{\partial \mathcal{G}}{\partial u}^T.$$

Thus, we are dealing with a second-order in time linear BVP problem.

It should also be noted that we are considering a specific form of the LSS sensitivity method without time dilation and its associated weighting parameters.[17] Although this may have an effect on the accuracy of the LSS adjoint, our proposed methods remain applicable to the LSS method with the time dilation term included.

B. Discretization of the LSS adjoint

1. *Primal-Problem Discretization*

The LSS method does not require a particular discretization method, but in order to make our explanation more concrete, we choose to demonstrate the LSS method using a Crank-Nicolson discretization. We discretize the IVP (1) using N uniform time steps; thus $\Delta t = T/N$ and $t_i = (i-1)\Delta t$, where $i = 1, 2, \ldots, N+1$. Including the initial condition, we have to solve

$$\begin{aligned} u_1 &= u_{\mathsf{IC}}, \\ R_i(\boldsymbol{u}, \nu, t) = \frac{u_{i+1} - u_i}{\Delta t} - \frac{1}{2}\left[\mathcal{F}(u_i, \nu, t_i) + \mathcal{F}(u_{i+1}, \nu, t_{i+1})\right] &= 0, \qquad \forall i = 1, 2, \ldots, N. \end{aligned} \tag{14}$$

In addition, we discretize the functional $\mathcal{J}$ using the midpoint rule[a]:

$$J(\boldsymbol{u}) = \sum_{i=1}^{N} \frac{1}{2}\left[\mathcal{G}(u_i, \nu, t_i) + \mathcal{G}(u_{i+1}, \nu, t_{i+1})\right] \Delta t, \tag{15}$$

where $\boldsymbol{u} = [u_1, u_2, \ldots, u_{N+1}]^T$ is the discrete solution.

2. *Standard Adjoint Discretization*

We can also use the midpoint rule for the term $\int_0^T \psi^T \mathcal{R}\, dt$ in the Lagrangian, see (5). Thus, the discretization of the Lagrangian becomes

$$\begin{aligned} L(u, \psi, \nu) = &\sum_{i=1}^{N} \frac{1}{2}\left[\mathcal{G}(u_i, \nu, t_i) + \mathcal{G}(u_{i+1}, \nu, t_{i+1})\right] \Delta t \\ &+ \sum_{i=1}^{N} \psi_{i+1}^T \left\{ \frac{u_{i+1} - u_i}{\Delta t} - \frac{1}{2}\left[\mathcal{F}(u_i, \nu, t_i) + \mathcal{F}(u_{i+1}, \nu, t_{i+1})\right] \right\} \Delta t + \psi_1^T (u_1 - u_{\mathsf{IC}}). \end{aligned}$$

We obtain the discrete adjoint equation by differentiating L with respect to the u_i and setting each partial derivative to zero. For $i = 1$ we get

$$\frac{\Delta t}{2} \left.\frac{\partial \mathcal{G}}{\partial u}\right|_1 - (\psi_2^T - \psi_1^T) - \psi_2^T \left[\frac{\Delta t}{2} \left.\frac{\partial \mathcal{F}}{\partial u}\right|_1\right] = 0.$$

Transposing and rearranging we have

$$\psi_1 = \psi_2 + \mathrm{O}(\Delta t),$$

which is a first-order accurate extrapolation of ψ_1 based on ψ_2.

Next, differentiating L with respect to u_{N+1} we get

$$\frac{\Delta t}{2} \left.\frac{\partial \mathcal{G}}{\partial u}\right|_{N+1} + \psi_{N+1}^T - \psi_{N+1}^T \left[\frac{\Delta t}{2} \left.\frac{\partial \mathcal{F}}{\partial u}\right|_{N+1}\right] = 0.$$

Transposing and rearranging gives

$$\psi_{N+1} = \frac{\Delta t}{2}\left[\psi_{N+1}^T \left.\frac{\partial \mathcal{F}}{\partial u}\right|_{N+1} - \left.\frac{\partial \mathcal{G}}{\partial u}\right|_{N+1}\right].$$

[a] It should be noted that when fully expanded, Equation (15) becomes the trapezoid rule.

As $\Delta t \to 0$, we recover $\psi_{N+1} = 0$. Thus, the above is a first-order approximation to the terminal condition in (7).

Finally, differentiating L with respect to u_i, where $i \neq 1$ and $i \neq N+1$, we get

$$\Delta t \left.\frac{\partial \mathcal{G}}{\partial u}\right|_i - \left(\psi_{i+1}^T - \psi_i^T\right) - \frac{\Delta t(\psi_{i+1}^T + \psi_i^T)}{2} \left.\frac{\partial \mathcal{F}}{\partial u}\right|_i = 0.$$

Transposing, dividing by Δt, and rearranging this equation, we get

$$-\frac{\psi_{i+1} - \psi_i}{\Delta t} - \left[\frac{\partial \mathcal{F}}{\partial u}\right]_1^T \frac{(\psi_{i+1} + \psi_i)}{2} + \left[\frac{\partial \mathcal{G}}{\partial u}\right]_i = 0,$$

which is the Crank-Nicholson discretization of the DE in (7).

Later, when discussing the discretization of the LSS adjoint, it will be helpful to have a compact representation of the discretization of the conventional discrete adjoint. To this end we write

$$\mathsf{A}^T \boldsymbol{\psi} + \boldsymbol{b} = \mathbf{0},$$

where

$$\boldsymbol{b} = \left[\frac{\Delta t}{2}\left[\frac{\partial \mathcal{G}}{\partial u}\right]_1, \Delta t\left[\frac{\partial \mathcal{G}}{\partial u}\right]_2, \ldots, \Delta t\left[\frac{\partial \mathcal{G}}{\partial u}\right]_N, \frac{\Delta t}{2}\left[\frac{\partial \mathcal{G}}{\partial u}\right]_{N+1}\right]^T,$$

and for the Jacobian we have

$$\mathsf{A} = \begin{bmatrix} \mathsf{I} & & & & & \\ \left[-\mathsf{I} - \frac{\Delta t}{2}\left.\frac{\partial \mathcal{F}}{\partial u}\right|_1\right] & \left[\mathsf{I} - \frac{\Delta t}{2}\left.\frac{\partial \mathcal{F}}{\partial u}\right|_2\right] & & & & \\ & \left[-\mathsf{I} - \frac{\Delta t}{2}\left.\frac{\partial \mathcal{F}}{\partial u}\right|_2\right] & \left[\mathsf{I} - \frac{\Delta t}{2}\left.\frac{\partial \mathcal{F}}{\partial u}\right|_3\right] & & & \\ & & \ddots & \ddots & & \\ & & & & \left[-\mathsf{I} - \frac{\Delta t}{2}\left.\frac{\partial \mathcal{F}}{\partial u}\right|_N\right] & \left[\mathsf{I} - \frac{\Delta t}{2}\left.\frac{\partial \mathcal{F}}{\partial u}\right|_{N+1}\right] \end{bmatrix}.$$

Once the adjoint $\boldsymbol{\psi}$ is obtained, a discrete version of (6) is used to obtain the gradient:

$$\sum_{i=1}^{N} \frac{1}{2}\left(\left.\frac{\partial \mathcal{G}}{\partial \nu}\right|_i + \left.\frac{\partial \mathcal{G}}{\partial \nu}\right|_{i+1}\right)\Delta t + \sum_{i=1}^{N} \psi_{i+1} \left.\frac{\partial R_i}{\partial \nu}\right|_i^T, \tag{16}$$

where R_i is defined in (14).

3. *Discretization of the LSS Adjoint*

Referring to (10), the optimization problem associated with the discretized LSS adjoint method is

$$\begin{aligned} \min_{\boldsymbol{u}'} \; & J(\bar{\boldsymbol{u}}) + \boldsymbol{b}^T \boldsymbol{u}' + \frac{1}{2}\left(\boldsymbol{u}'\right)^T \mathsf{H}\boldsymbol{u}' \\ \text{s.t.} \quad & \mathsf{A}'\boldsymbol{u}' = \mathbf{0}. \end{aligned} \tag{17}$$

In Equation (17), $J(\bar{\boldsymbol{u}}) + \boldsymbol{b}^T\boldsymbol{u}'$ is the discrete version of the linearized objective function; here, $\bar{\boldsymbol{u}}$ is the discretized reference state and $\boldsymbol{u}'$ is the discretized perturbation. The term $\frac{1}{2}\left(\boldsymbol{u}'\right)^T \mathsf{H}\boldsymbol{u}'$ is the discretized regularization term, where $\mathsf{H} = \Delta t\,\mathsf{diag}\,(1/2, 1, 1, \ldots, 1, 1/2)$ holds the quadrature weights corresponding to

the trapezoid rule. Again, since the LSS method removes the initial condition from the linearized problem, we remove the first block row from A to give A' in (17). To be explicit, A' is defined as

$$\mathsf{A}' = \begin{bmatrix} \left[-\mathsf{I} - \frac{\Delta t}{2}\left.\frac{\partial \mathcal{F}}{\partial u}\right|_1\right] & \left[\mathsf{I} - \frac{\Delta t}{2}\left.\frac{\partial \mathcal{F}}{\partial u}\right|_2\right] & & \\ & \left[-\mathsf{I} - \frac{\Delta t}{2}\left.\frac{\partial \mathcal{F}}{\partial u}\right|_2\right] & \left[\mathsf{I} - \frac{\Delta t}{2}\left.\frac{\partial \mathcal{F}}{\partial u}\right|_3\right] & \\ & & \ddots & \ddots \\ & & & \left[-\mathsf{I} - \frac{\Delta t}{2}\left.\frac{\partial \mathcal{F}}{\partial u}\right|_N\right] & \left[\mathsf{I} - \frac{\Delta t}{2}\left.\frac{\partial \mathcal{F}}{\partial u}\right|_{N+1}\right] \end{bmatrix}$$

The optimization problem (17) is a quadratic optimization problem with a convex objective. The solution to such an optimization problem is given by the solution of the following saddle-point linear system:

$$\begin{bmatrix} \mathsf{H} & (\mathsf{A}')^T \\ \mathsf{A}' & 0 \end{bmatrix} \begin{bmatrix} \boldsymbol{u}' \\ \boldsymbol{\psi} \end{bmatrix} = \begin{bmatrix} -\boldsymbol{b} \\ \mathbf{0} \end{bmatrix}$$

However, since $\mathsf{A}'\boldsymbol{u}' = 0$ (the second row of the above), we can multiply the first block row by $\mathsf{A}'\mathsf{H}^{-1}$ to get a single equation for the discrete LSS:

$$\mathsf{A}'\mathsf{H}^{-1}\left(\mathsf{A}'\right)^T \boldsymbol{\psi} = -\mathsf{A}'\mathsf{H}^{-1}\boldsymbol{b}. \tag{18}$$

This is effectively forming the Schur complement of the block system. With the LSS adjoint found, (16) is used in the same way as the conventional adjoint to obtain the desired derivative.

C. A Petrov-Galerkin approach to the LSS method: Fourier approximation of the adjoint

Typically the adjoint and the state spaces are the same, which simplifies the mathematical formulation. This is the case for the standard LSS adjoint, but the computational cost challenges with the LSS problem remain. Motivated by this, we seek to reduce the dimension of the discretized LSS system.

Our general approach is to use a Petrov-Galerkin approach for the linearized state equation in (10). That is, we select different solution spaces for the state and LSS adjoint. In particular, and unlike typical finite-element discretizations, the dimension of the state and adjoint spaces can be different.

To illustrate the more general idea, we approximate the LSS adjoint in the Lorenz system using a truncated Fourier sine series. In the discrete case, this is done by projecting the LSS Jacobian A' (which, again, is A with the first block row removed) onto a series of sine terms. In matrix form, this projection is defined by:

$$\mathsf{A}_{\text{sin}} = \mathsf{F}^T\mathsf{A}', \tag{19}$$

where

$$\mathsf{F}^T = \begin{bmatrix} \sin\left(\pi\frac{t_1}{T}\right) & \sin\left(\pi\frac{t_2}{T}\right) & \dots & \sin\left(\pi\frac{t_N}{T}\right) \\ \sin\left(2\pi\frac{t_1}{T}\right) & \sin\left(2\pi\frac{t_2}{T}\right) & \dots & \sin\left(2\pi\frac{t_N}{T}\right) \\ \vdots & \vdots & & \vdots \\ \sin\left(M\pi\frac{t_1}{T}\right) & \sin\left(M\pi\frac{t_2}{T}\right) & \dots & \sin\left(M\pi\frac{t_N}{T}\right) \end{bmatrix}, \tag{20}$$

where M is the number of Fourier modes. The approximated Jacobian is then used to compute the adjoint vector in a similar manner to the original LSS method; see (18):

$$\left(\mathsf{F}^T\mathsf{A}'\right)\mathsf{H}^{-1}\left(\mathsf{F}^T\mathsf{A}'\right)^T\boldsymbol{\psi}_{\text{sin}} = -\left(\mathsf{F}^T\mathsf{A}'\right)\mathsf{H}^{-1}\boldsymbol{b}, \tag{21}$$

which can then be used to obtain values for the gradient. The process of computing the gradient from the adjoint is similar to that in (16), but involves the basis function values to reconstitute the adjoint using

$$\psi_i = \sum_{m=1}^{M} \psi_{\text{sin}} \sin\left(m\pi t_i\right), \tag{22}$$

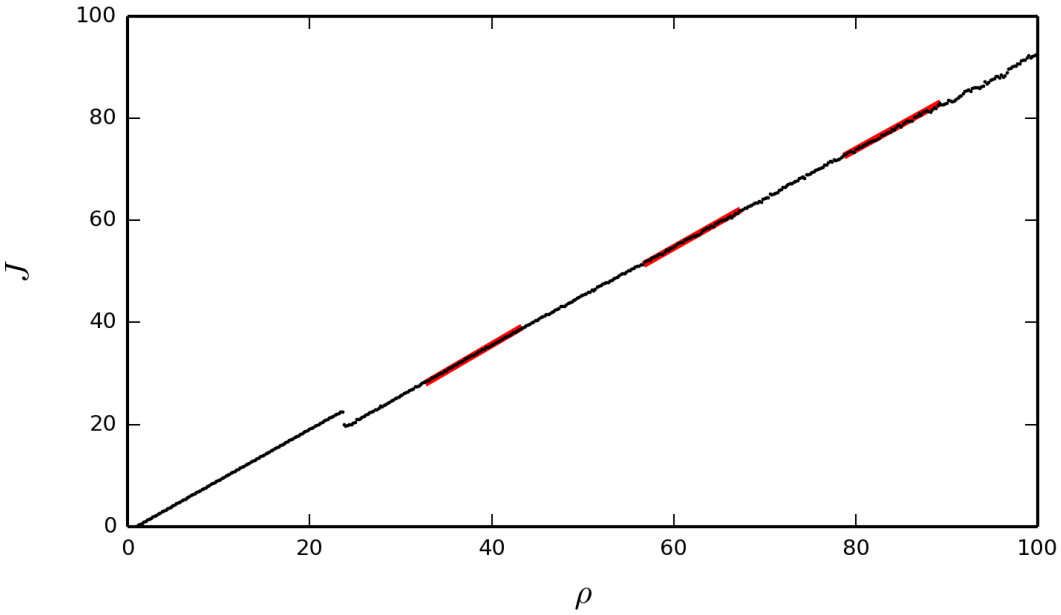

Figure 3. A plot of the objective function $\mathcal{J} = 1/T \int_0^T z$ versus ρ. Overlaid in red are some examples of the LSS-calculated gradient; it is apparent that there is good agreement between the gradient calculation and the sampled data.

where i is the time step, t_i is the simulation time at time step i, and M is the number of modes.

The process of solving Equation (18) for the adjoint is the computationally intensive aspect of the LSS method, as it requires a matrix solve on the typically quite large A. It is our hypothesis that the Fourier system will be at least an order of magnitude smaller than the standard system, and thus the computation of the adjoint using the Fourier-approximated method (21) has the potential to offer significant computational advantages.

IV. Numerical results

In order to compare the Fourier LSS method with the standard LSS method, both adjoints were computed and compared for simulations of the Lorenz model problem. In all simulations, the integration period was $T = 16$ and the time step was $\Delta t = 0.004$. Additionally, the solution of the Lorenz system was allowed to undergo a spin-up period of $T = 10$, to ensure that the gathered statistics better reflect the long-time behavior of the system. Typical values for the Lorenz system parameters were used: $\sigma = 10$, $\beta = 8/3$, and $\rho = 28$ (except where stated otherwise).

A. The standard LSS applied to the Lorenz problem

There are several factors that must be studied in order to assess the quality of the Fourier approximation. But first, it is important to quantify the accuracy of the standard LSS method, since this is the benchmark method. Figure 3 shows a representative set of three LSS-calculated gradients appearing together with a plot of the actual objective function (2) versus ρ. It is important to reiterate that with a standard adjoint method, the gradient lines (as can be seen in Figure 2) are nearly vertical, as they become unbounded due to the aforementioned instability of chaotic adjoint problems.

Figure 4 shows a sample of LSS-based gradients for the objective function with respect to ρ, over a range of $[0, 100]$. At each value of ρ, a Lorenz simulation was started with an initial condition of $[x, y, z] = [\Delta x, \Delta y, \rho + \Delta z]$, where Δx, Δy, and Δz are normally distributed random variables with a mean of zero and a standard deviation of 0.05. After setting the initial condition and solving the primal problem, the LSS adjoint was computed and used to obtain gradient information. An ensemble of twenty simulations was

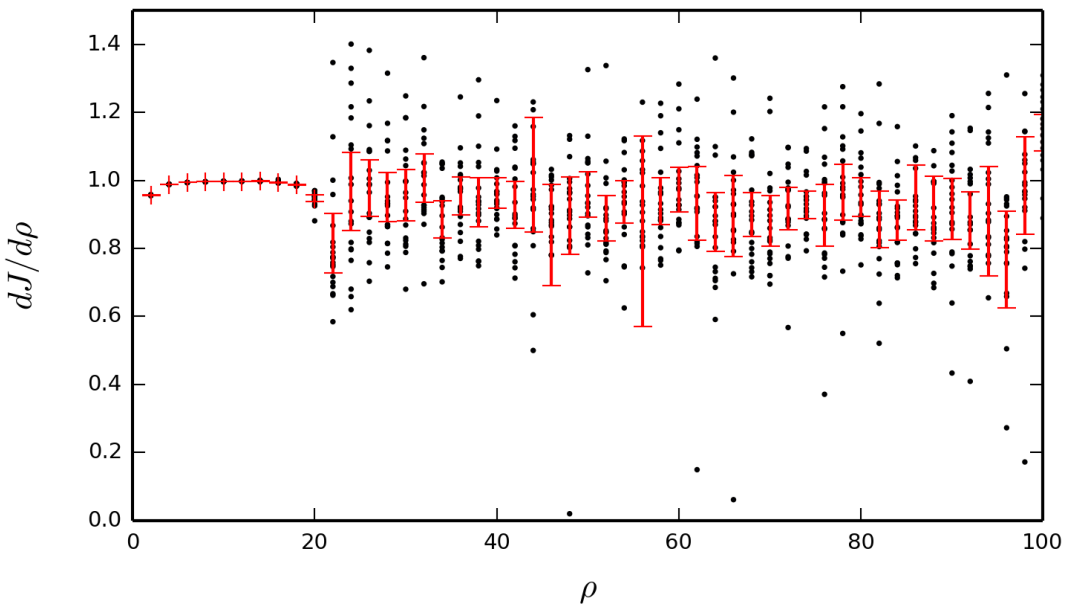

Figure 4. A statistical gathering of the gradient of the objective function with respect to ρ, calculated using the standard LSS adjoint. Each dot represents one simulation of the Lorenz system with $T = 16$; there are 20 simulations perfomed at each value of ρ, with slight randomness given to the initial conditions. Error bars, centered about the mean and with half-length one standard deviation, are shown. The locations in which the error bars appear longer than the dataset contain negative gradients which affect the calculation of the standard deviation.

conducted at each value of ρ. While the sensitivity to changes in initial conditions can be seen, the figure clearly shows the clustering of gradient values around $d\mathcal{J}/d\rho = 0.96$, which is the approximate value of the gradient for the chosen objective.[13] Recall that our LSS implementation does not include the time-dilation term, and we believe this explains the increased spread in the gradient values relative to the literature. The time-dilation term will be included in future work.

B. Fourier-approximated LSS

With the implementation of the standard LSS adjoint verified, we can begin investigating the quality of our Fourier approximation method. In all of the following results, the number of Fourier modes was selected as $M = 460$, which struck a balance between minimizing the relative gradient error and keeping the dimensions of the dense matrix suitably small. A qualitative comparison can be seen in Figure 5, in which the LSS-based adjoint and the Fourier-approximated adjoint are plotted in phase space. It is evident from Figure 5 that the approximation mirrors the standard LSS quite well; this is also observed by the more quantitative comparison of Figure 6, in which the difference in the x-coordinate of the adjoints over time is shown.

The same aggregate analysis shown in Figure 4 was conducted for the Fourier-approximated LSS, and these results can be seen in Figure 7. Again, the clustering of gradient values about $d\mathcal{J}/d\rho = 1.0$ is seen. The dip in the mean gradient at high values of ρ, compared to the standard LSS gradient sampling, is due to an insufficient number of modes to approximate the dominant dynamics of the problem. It should be noted that as ρ changes significantly, the number of modes required to drive the relative gradient error down tends to increase. A separate numerical experiment, using 800 modes instead of the 460 modes in Figure 7, confirmed that high values of ρ require more modes to approximate the LSS adjoint. Additional work is necessary to determine an automatic way of selecting the appropriate number of modes in the approximation.

Another verification of the Fourier approximation was conducted by analyzing the relative gradient error between the standard LSS and the Fourier-approximated LSS, defined in Equation (23), as a function of the

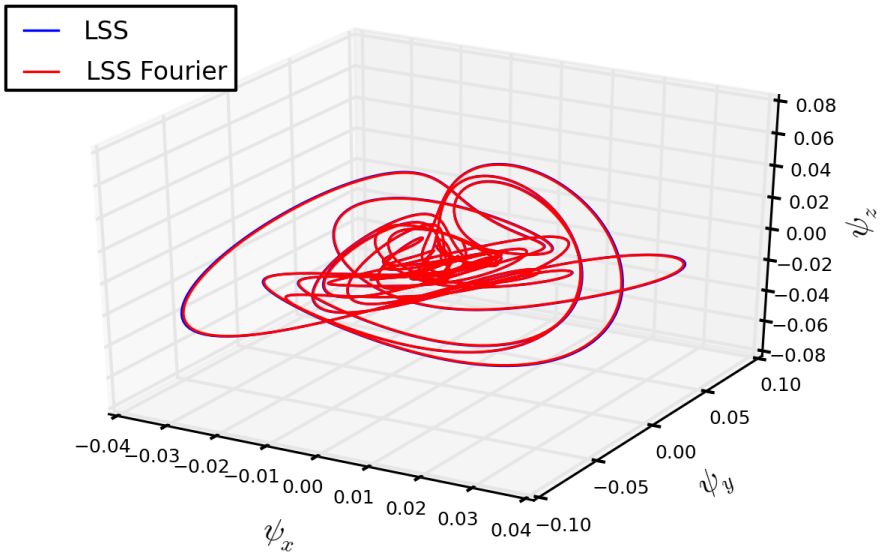

Figure 5. The standard LSS adjoint plotted in state space, alongside its Fourier approximation using 460 modes and a solution time of $T = 16$. Qualitatively, the closeness of the approximation can be seen.

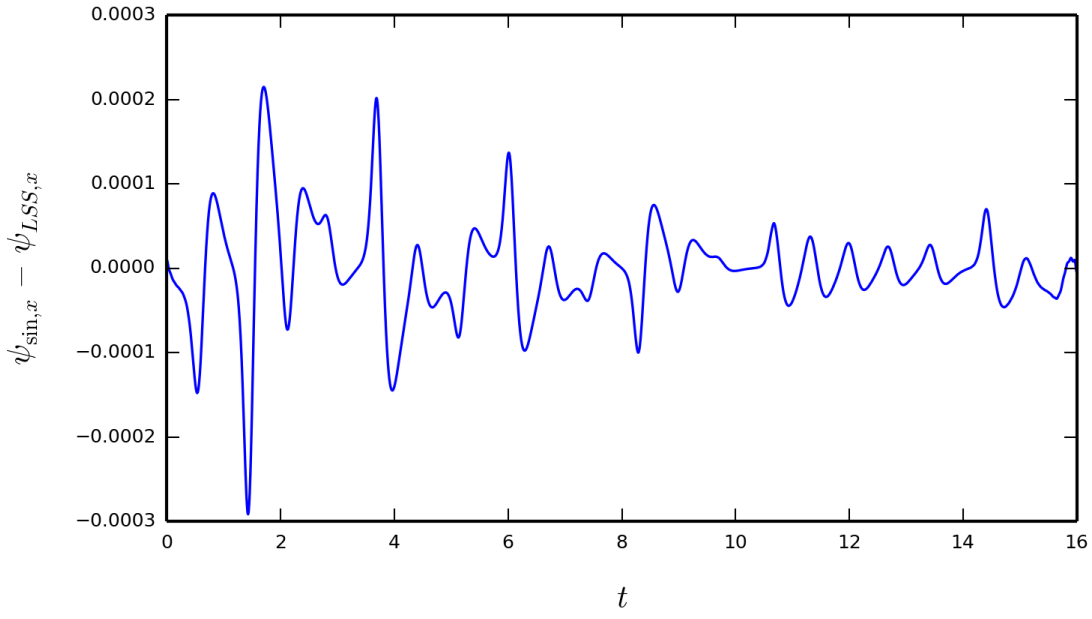

Figure 6. The error between the LSS adjoint and the Fourier-approximated LSS adjoint for the component of ψ_x corresponding to the x state variable plotted versus time.

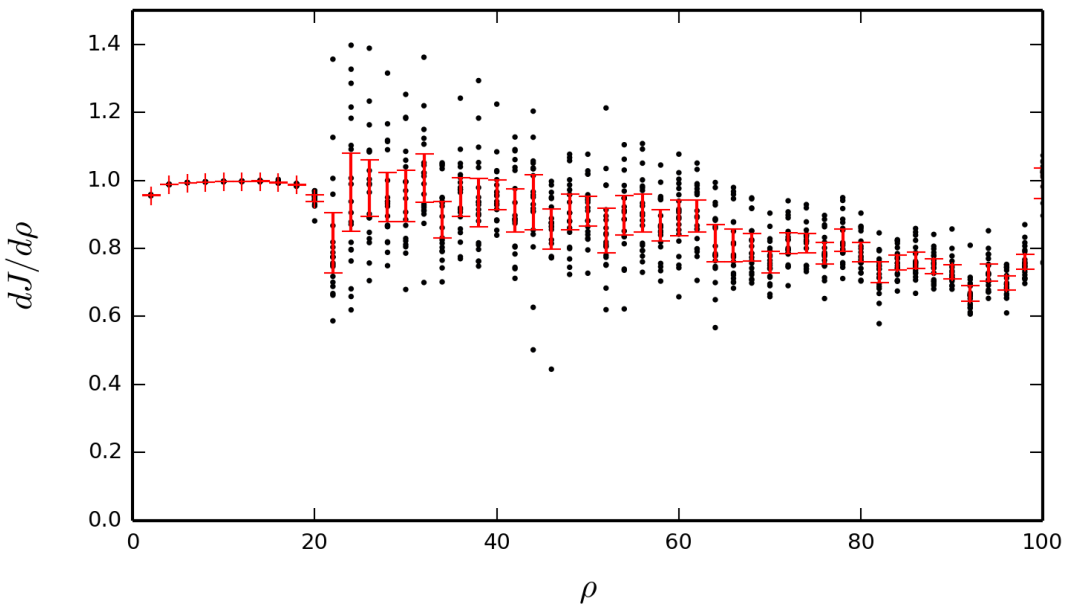

Figure 7. A statistical gathering of the gradient of the objective function with respect to ρ, calculated using the Fourier approximation to the LSS adjoint. Each dot represents one simulation of the Lorenz system with $T = 16$; there are 20 simulations perfomed at each value of ρ, with slight randomness given to the initial conditions. Error bars, centered about the mean and with half-length one standard deviation, are shown.

number of sine modes.

$$\epsilon_{dJ/d\rho} = \frac{\left|\left(\frac{dJ}{d\rho}\right)_{\text{LSS}} - \left(\frac{dJ}{d\rho}\right)_{\text{Fourier}}\right|}{\left|\left(\frac{dJ}{d\rho}\right)_{\text{LSS}}\right|}. \tag{23}$$

As can be seen in Figure 8, the relative gradient error decreases significantly once a sufficient number of modes are used, eventually plateauing on the order of 10^{-3}. This investigation was conducted at a parameter value of $\rho = 28$.

One disadvantage of the Fourier approximation is the condition number of the Schur complement matrix in (21). The condition number of the Schur complement is plotted against the number of modes in Figure 9 for $\rho = 28$. The condition number becomes quite high as the number of modes is increased, which has a negative effect on accuracy and preconditioning for iterative methods. Therefore, it is desirable to select a number of modes that captures desired level of accuracy, but does not overshoot the error plateau too far. However, some overshoot does not impact the condition number significantly, as the condition number grows rather slowly after the plateau.

A brief discussion of the 'plateau' in Figures 8 and 9 is warranted. At around 460 modes for this selection of ρ, the relative gradient error of the Fourier approximation stops decreasing. The main reason for this is that there will be some error inherent to the standard LSS method, due to the choice in discretization, and that the Fourier approximation is simply trying to capture high-frequency, broadband error. This is supported by the fact that smaller time step values in the discretization drives downward the resting value of the plateau that the relative gradient error reaches, visible in Figure 8.

Although this approximation method displays promise when compared to the standard LSS approach, it must be stated that the Fourier system matrix is dense. This could be a significant drawback, and future work is needed to investigate strategies to solve this system in the context of large-scale simulations.

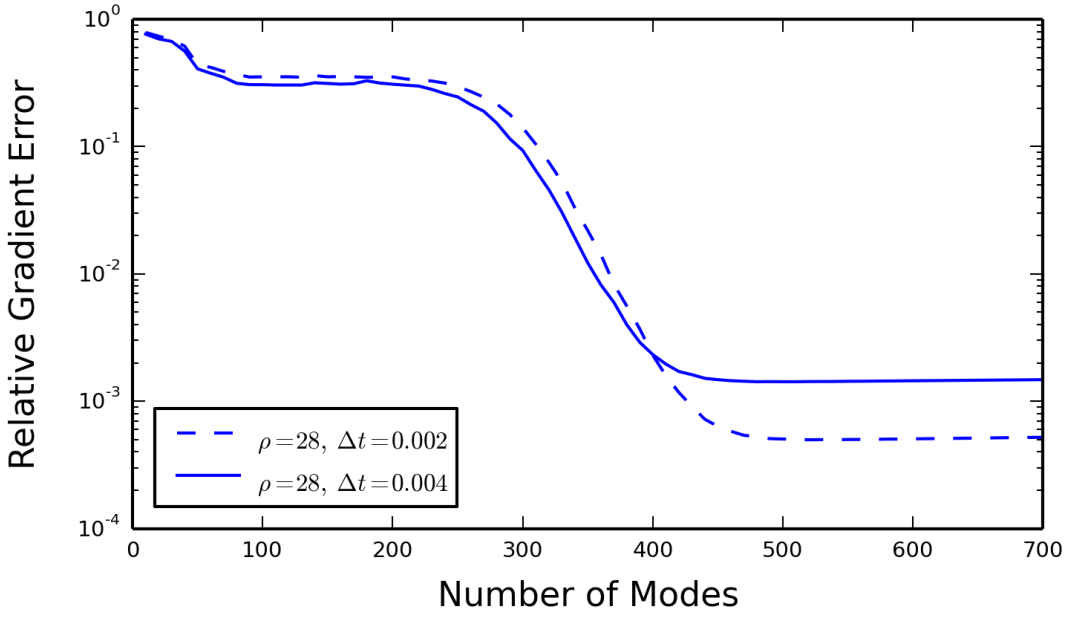

Figure 8. Relative error in gradient in the Fourier approximation of the least-squares shadowing adjoint matrix as the number of Fourier modes is varied, for differing values of Δt. Simulation time remains $T = 16$.

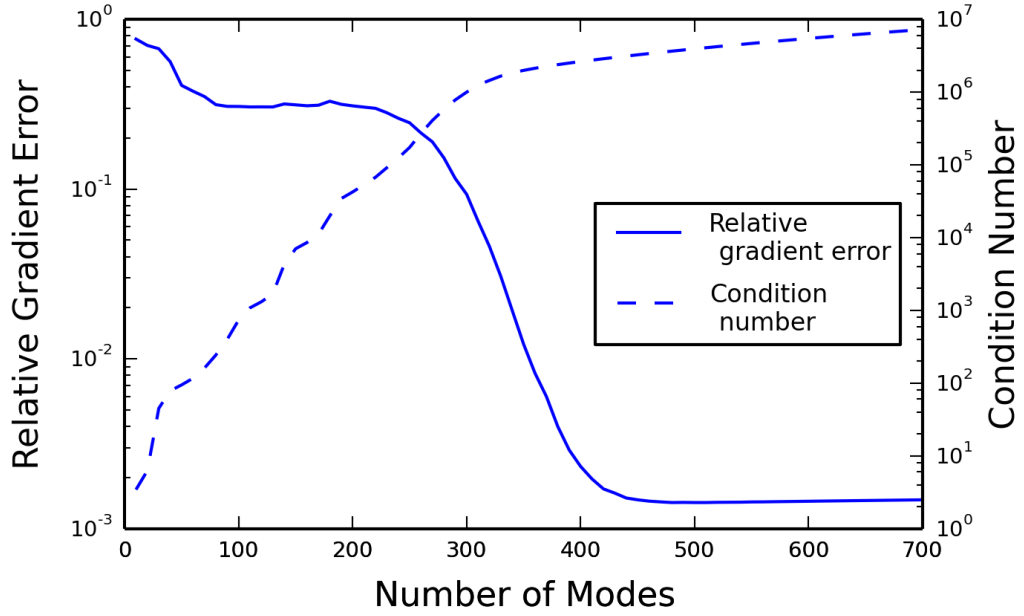

Figure 9. Relative error in gradient (blue) and condition number of the approximation of the least-squares shadowing adjoint matrix (red) as the number of Fourier modes is varied. Simulation time $T = 16$, and $\rho = 28$.

V. Conclusions

Chaotic problems are prevalent in many engineering disciplines, but their optimization with gradient-based methods is challenging, because the linear adjoint is unstable and grows without bound. The least-squares shadowing (LSS) sensitivity method has been proposed to obtain gradient information from chaotic problems in a stable fashion. While this method provides useful gradient information that is of reasonable accuracy, it is prohibitively computationally expensive for large-scale simulations. Motivated to reduce this cost, we have proposed a method of representing the adjoint solution in a discrete space that differs from the state space. In our model problem, we approximated the adjoint of the LSS method using a Fourier series. We then conducted several numerical experiments to investigate the quality of our approximation. The relative gradient error remains bounded, and the condition number of the least-squares adjoint remains tractable. The idea of using distinct spaces for the state and LSS adjoint appears to merit further investigation and application to other, more difficult, chaotic problems.

VI. Acknowledgments

A. Ashley and J. Hicken gratefully acknowledge the support of the Air Force Office of Scientific Research Award FA9550-15-1-0242 under Dr. Jean-Luc Cambier. We also thank RPI's Scientific Computation Research Center for the use of computer facilities.

References

[1]Raveh, D. E., Levy, Y., and Karpel, M., "Structural optimization using computational aerodynamics," *AIAA Journal*, Vol. 38, No. 10, 2000, pp. 1974-1982.

[2]Jameson, A., "Aerodynamic design via control theory," *Journal of Scientific Computing*, Vol. 3, No. 3, 1988, pp. 233-260.

[3]Pironneau, O., "On optimum design in fluid mechanics," *Journal of Fluid Mechanics*, Vol. 64, Part 1, 1974, pp. 97-110.

[4]Nemec, M., and Zingg, D. W., "Multipoint and multi-objective aerodynamic shape optimization," *AIAA Journal*, Vol. 42, No. 6, 2004, pp. 1057-1065.

[5]Anderson, W. K., and Bonhaus, D. L., "Airfoil design on unstructured grids for turbulent flows", *AIAA Journal*, Vol. 37, Feb. 1999, pp. 185-191.

[6]Nielsen, E. J., and Anderson, W. K., "Aerodynamic design optimization on unstructured meshes using the Navier-Stokes equations", *AIAA Journal*, Vol. 37, Nov. 1999, pp. 1411-1499.

[7]Reuther, J. J., Jameson, A., Alonso, J. J., Rimlinger, M. J., and Saunders, D., "Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers, part 1", *AIAA Journal*, Vol. 36, Jan. 1999, pp. 51-60.

[8]Reuther, J. J., Jameson, A., Alonso, J. J., Rimlinger, M. J., and Saunders, D., "Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers, part 2", *AIAA Journal*, Vol. 36, Jan. 1999, pp. 61-74.

[9]Maute, K., Nikbay, M., Farhat, C., "Sensitivity analysis and design optimization of three-dimensional non-linear aeroelastic systems by the adjoint method", *International Journal for Numerical Methods in Engineering*, Vol. 56, Feb. 2003, pp.911-933.

[10]Kenway, G. K. W. and Martins, J. R. R. A., "Multipoint high-fidelity aerostructural optimization of a transport aircraft configuration", *Journal of Aircraft*, Vol. 51, 2014, pp. 144-160.

[11]Kenway, G. K. W., Kennedy, G. J., and Martins, J. R. R. A., "Scalable parallel approach for high-fidelity steady-state aeroelastic analysis and adjoint derivative computations", *AIAA Journal*, Vol. 52, 2014, pp. 935-951.

[12]Martins, J. R. R. A., Alonso, J. J., and Reuther, J. J., "High-fidelity aerostructural design optimization of a supersonic business jet," *Journal of Aircraft*, Vol. 41, No. 3, 2004, pp. 523-530.

[13]Lea, D.J., Allen, M. R., and Haine, T. W. N., "Sensitivity analysis of the climate of a chaotic system," *Tellus A*, Vol. 52, No. 5, Oct. 2000, pp. 523-532.

[14]Eyink, G. L., Haine, T. W. N., and Lea, D. J., "Ruelle's linear response formula, ensemble adjoint schemes, and Levy flights," *Nonlinearity*, Vol. 17, No. 5, Sep. 2004, pp. 1867+.

[15]Ashley, A. and Hicken, J. E., "Optimization algorithm for systems governed by chaotic dynamics," *2014 AIAA Aviation Conference*, Jun. 2014, AIAA 2014-2434.

[16]Wang, Q., Hu, R., Blonigan, P., "Least Squares Shadowing sensitivity analysis of chaotic limit cycle oscillations," *Journal of Computational Physics*, Jun. 2014, 267:210-224.

[17]Liao, H., "Efficient sensitivity analysis method for chaotic dynamical systems," *Journal of Computational Physics*, Vol. 313, Feb. 2016, pp. 57-75.