Reducing the cost of shadowing-based adjoint sensitivity analysis for turbulent flows

By P. J. Blonigan†, S. M. Murman‡ AND A. Towne

Adjoint-based sensitivity analysis methods are powerful tools for engineers who use flow simulations for design. However, the conventional adjoint method breaks down for scale-resolving simulations such as large-eddy simulation (LES) or direct numerical simulation (DNS), which exhibit the chaotic dynamics inherent in turbulent flows. Sensitivity analysis based on least-squares shadowing (LSS) avoids the issues encountered by conventional methods, but has a high computational cost. This report analyzes some properties of one approach, non-intrusive LSS (NILSS), and discusses the implications for computational cost reductions of the approach. We find that there are opportunities to considerably reduce the cost of NILSS, but simple linear dimension reduction will not suffice.

1. Introduction

Scale-resolving simulations such as LES are necessary for engineering design and flow analysis, most notably flows in which jets, wakes, and separation dominate. In these cases, the Reynolds-averaged Navier-Stokes (RANS) solvers typically used by today's engineers often fail to accurately capture the relevant flow physics (Leonard *et al.* 2015). At the same time, engineers are interested in gradient-based design optimization, error estimation, and uncertainty quantification with flow simulations. All of these require efficient approaches for sensitivity analysis. Unfortunately, conventional sensitivity analysis approaches such as the adjoint method do not compute accurate sensitivities for statistically stationary quantities of interest in scale-resolving turbulent flow simulations such as LES or DNS (Blonigan *et al.* 2018). This is because, unlike RANS, LES and DNS resolve the chaotic dynamics of turbulent fluid flows (Keefe *et al.* 1992), and the adjoint method computes unusable sensitivities for chaotic systems (Lea *et al.* 2000).

Shadowing-based sensitivitiy analysis has shown great promise for computing accurate sensitivities of statistically stationary quantities in chaotic dynamical systems. These approaches are based off on LSS introduced in Wang *et al.* (2014). LSS has been shown to compute accurate sensitivities for a number of chaotic dynamic systems, including a two-dimensional airfoil with chaotic vortex shedding (Blonigan *et al.* 2018). However, this study also demonstrated the excessive computational costs of LSS. Recently, NILSS was used to compute sensitivities for chaotic flow around a cylinder (Ni & Wang 2017) and the minimal flow unit, a channel flow with a truncated domain (Blonigan 2017). In the latter study, NILSS was used to compute the adjoint field for the minimal flow unit, and the adjoint field provided insight into the near-wall flow structures and turbulent bursting events (Blonigan 2017). However, both studies demonstrated the high cost of NILSS, which scales with the number of positive Lyapunov exponents present in a simulation. This number is most likely very large for flows of practical interest (Keefe *et al.* 1992; Blonigan *et al.* 2016).

† NASA Postdoctoral Program, Universities Space Research Association
‡ NASA Ames Research Center

This report analyzes the potential for cost reductions of NILSS. Specifically, the ensemble of tangent equations is analyzed, as this part of the NILSS algorithm contributes most of the computational cost.

2. Breakdown of conventional sensitivity analysis

Consider a flow simulation with the following nonlinear governing equations

$$\frac{du}{dt} = f(u; s) \tag{2.1}$$

where u is a length n vector of state variables and s is some system parameter. For a three-dimensional compressible flow simulation, u contains the five conserved quantities at all degrees of freedom (grid points). The parameter s could be a flow parameter such as the freestream Mach number or a geometric parameter such as chord length.

When designing a system with unsteady flow, engineers are often interested in a time-averaged quantity $\bar{J}$,

$$\bar{J}(s) = \frac{1}{T}\int_{t_0}^{t_0+T} J[u(t; s); s]\ dt, \tag{2.2}$$

where $J[u(t; s); s]$ is some instantaneous quantity of interest, such as the lift or drag on an airfoil. In many cases, including applications with turbulent flow, engineers are interested in infinite time averages, $\bar{J}$ as $T \to \infty$. Since the exact evaluation of this is not computationally feasible, the infinite time average is approximated with a choice of T that ensures $\bar{J}(s)$ is nearly stationary (i.e. does not vary with T) (Oliver *et al.* 2014).

Sensitivities with respect to the parameter s can be computed using the following equation obtained by differentiating Eq. (2.2)

$$\frac{d\bar{J}}{ds} = \frac{1}{T}\int_{t_0}^{t_0+T} \left(\left\langle \frac{\partial J}{\partial u}, v \right\rangle + \frac{\partial J}{\partial s} \right)\ dt, \qquad v \equiv \frac{\partial u}{\partial s}, \tag{2.3}$$

where $\langle \cdot, \cdot \rangle$ is the inner product and all variables on the right-hand side are time dependent.

Conventionally, the tangent solution, v, is obtained from the linearization of Eq. (2.1), referred to as the tangent equation

$$\frac{dv}{dt} = \frac{\partial f}{\partial u} v + \frac{\partial f}{\partial s}, \qquad v(t_0) = \frac{\partial u_0}{\partial s} = 0. \tag{2.4}$$

The conventional approach using Eqs. (2.3) and (2.4) to compute sensitivities works for periodic and quasi-periodic systems if the time horizon is an integer number of periods or if windowing is used. However, it fails for chaotic dynamical systems because chaotic systems have at least one positive Lyapunov exponent. That is, the unsteady tangent equations for a chaotic dynamical system will grow exponentially in time when exposed to certain perturbations. Although the initial condition for the tangent equation, Eq. (2.4), is zero, the term $\partial f/\partial s$ acts like a forcing term, which almost always contains components that align with the unstable subspace, leading to exponential growth of the tangent solution.

The exponential growth of the magnitude of $v(t)$ means that as the time horizon length T is increased, the gradient computed using Eq. (2.3) grows exponentially as well. This means that conventional sensitivity analysis will compute very large, unusable

sensitivities for chaotic dynamical systems such as scale-resolving flow simulations including LES. This result was explained using forward sensitivity analysis, but conventional adjoint-based sensitivity analysis encounters the same exponentially growth backward in time.

3. Non-intrusive least-squares shadowing

One approach to avoid the breakdown discussed in the previous section is LSS (Wang *et al.* 2014). Although LSS can compute accurate sensitivities, these sensitivities come at a relatively high cost, most of which comes from solving a minimization problem that scales with the product of spatial degrees of freedom and the number of discrete time steps within the time horizon of interest, $t_0 < t < t_0 + T$. NILSS reduces the size of this minimization problem considerably. Specifically, in NILSS the following alternative minimization problem is solved

$$\min_{v(t_i)} \sum_{i=0}^{K} \|v(t_i)\|_2^2, \qquad \text{s.t.} \quad \frac{dv}{dt} = \frac{\partial f}{\partial u} v + \frac{\partial f}{\partial s} + \eta f, \quad t \in [t_0, t_K], \tag{3.1}$$

where η is chosen so that $\langle v(t), f[u(t); s] \rangle = 0$. Now the tangent solution norm $\|v(t)\|_2$ is minimized at $K + 1$ checkpoints t_i instead of at all time steps between t_0 and t_K. In this case, the minimization problem, Eq. (3.1), can be solved with a $Kn \times Kn$ Karush-Kuhn-Tucker (KKT) system.

The size of the minimization problem can be further reduced by decomposing the tangent solution $v(t)$ into a weighted combination of one forced and p unforced components, $\hat{v}(t)$ and $W^j(t)$, respectively,

$$v(t) = v_i^*(t) + W_i(t)\alpha_i, \quad t_{i-1} \le t < t_i, \tag{3.2}$$

$$\frac{dv_i^*}{dt} = \frac{\partial f}{\partial u} v_i^* + \frac{\partial f}{\partial s} + \eta^* f, \quad t_{i-1} \le t < t_i, \tag{3.3}$$

$$\frac{dW_i^j}{dt} = \frac{\partial f}{\partial u} W_i^j + \eta^j f, \quad t_{i-1} \le t < t_i, \tag{3.4}$$

where α_i is a length-p vector, W_i is a $n \times p$ matrix, and W_i^j is the jth column of W_i.

This decomposition allows the minimization statement, Eq. (3.1), to be written as a minimization over the weights α_i at each checkpoint t_i rather than the entire tangent solution $v(t_i)$. Since exponentially growing components of $v(t)$ will contribute the most by far to $\|v(t_i)\|_2^2$, Eq. (3.1) can be approximated by setting $W_i(t_i)$ equal to a set of vectors spanning the unstable subspace at t_i. This choice of $W_i(t_i)$ reduces the size of the KKT system to $Kp \times Kp$.

Setting $W_i(t_i)$ to span the unstable subspace allows α_i to be chosen such that the solutions of Eq. (3.4) can cancel out the exponential growth in the solution $\hat{v}(t)$ of Eq. (3.3), resulting in no exponential growth in $v(t)$. This means that the number of unforced tangents p should be at least the number of positive Lyapunov exponents, n^+. Past studies, including those by Pulliam & Vastano (1993), Keefe *et al.* (1992), and Sirovich & Deane (1991), have found n^+ to be a small fraction of n for a range of different flows.

The implementation of adjoint NILSS is an algorithm comprised of three main parts. These parts are summarized below, and Appendix A contains the entire NILSS algorithm. Readers are referred to Blonigan (2017) for the derivation and further details on the approach.

In the first part of the NILSS algorithm, at least n^+ tangent solutions are solved over the time horizon of interest. A QR decomposition is computed of the $n \times n^+$ matrices $W_i(t_i)$ at all K checkpoints. In the second part, the R matrices from these decompositions, along with the gradient contribution g_i from each segment (see appendix A) are used to form the NILSS KKT linear system.

$$\begin{bmatrix} R_1 R_1^T + I & -R_2 & & & \\ -R_2^T & R_2 R_2^T + I & -R_3 & & \\ & \ddots & \ddots & \ddots & \\ & -R_{K-1}^T & R_{K-1} R_{K-1}^T + I & -R_{K-1} & \\ & & -R_K^T & R_K R_K^T + I \end{bmatrix} \begin{pmatrix} \psi_1 \\ \psi_2 \\ \vdots \\ \psi_K \end{pmatrix} = \begin{pmatrix} R_1 g_1 - g_2 \\ R_2 g_2 - g_3 \\ \vdots \\ R_K g_K \end{pmatrix}. \tag{3.5}$$

The third and final part of the algorithm involves solving the following adjoint equation over all K time segments in reverse from K to 1.

$$-\frac{d\hat{w}}{dt} = \frac{df}{du}\bigg|_t^T \hat{w} + \frac{1}{\Delta T}\frac{\partial J}{\partial u}\bigg|_t. \tag{3.6}$$

The terminal condition of the adjoint in each segment depends on the next segment and the solution of the NILSS KKT system (3.5). Finally, the sensitivity can be computed using the conventional adjoint sensitivity equation

$$\frac{d\bar{J}}{ds} = \int_{t_0}^{t_K} \frac{\partial f}{\partial s}\bigg|_t^T \hat{w}(t)\, dt + \frac{\partial \bar{J}}{\partial s}. \tag{3.7}$$

In practice, the vast majority of the cost of NILSS is from step 5 in the first algorithm, A, solving the tangent equations (Blonigan 2017). Therefore, the cost of NILSS for a given flow simulation scales roughly with the number of positive Lyapunov exponents, n^+. Although n^+ is much smaller than n, it is still relatively large for scale-resolving turbulent flow simulations. For example, Blonigan *et al.* (2016) estimated n^+ to be roughly 1500 for $Re_\tau = 180$ channel flow in a $4\pi \times 2 \times 2\pi$ domain. This means that NILSS would cost roughly 1500 times as much as a primal solution for a given time horizon.

4. Tangent equation ensemble cost

Since the cost of NILSS scales with the cost of the ensemble of tangent equations, reducing the cost of this ensemble is the most effective way to reduce the overall cost of NILSS. To achieve this cost reduction, further understanding of the tangent equations and the unstable subspace for turbulent flows is required.

4.1. *Tangent solution accuracy requirements*

One way to reduce computational costs is to approximately compute the ensemble of tangent equations. For instance, if one were solving the tangent equations with implicit time stepping, the residual tolerance for the tangent could be made lower than that for the primal equations. One could also consider using reduced-order models of the tangent equation, which can be very inexpensive but have some error associated with them.

To study how accurately the tangent equations should be solved, a simple error model is added to the discrete homogeneous tangent equation

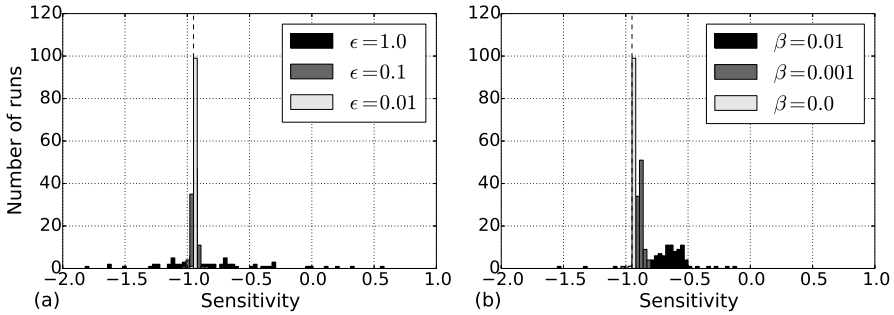

FIGURE 1. Distribution of the sensitivity $d\bar{J}/dc$ for a fixed primal solution $u(x,t)$ but different values of ϵ (a) and β (b). Note that $\beta = 0.0$ in (a) and $\epsilon = 0.01$ in (b). Dotted lines indicate the sensitivity computed by NILSS with the true tangent equations. All cases were run with a $T = 500$ time horizon and $K = 50$ segments. 100 sample runs with random Q_0 were performed for each ϵ, β pair.

$$\frac{dW_i^{'j}}{dt} = \left(\frac{\partial f}{\partial u} + \epsilon S\right) W_i^{'j} + \beta, \quad t_{i-1} \leq t < t_i, \tag{4.1}$$

where S is a diagonal matrix whose main diagonal entries are random numbers sampled from a uniform distribution $\mathcal{U}(-1,1)$, ϵ is the amplitude of the error term, and β is a bias term. This error model ensures that the error amplitude remains proportional to the tangent, which is important since the tangent grows exponentially for the unstable directions, quickly rendering a fixed amplitude error forcing term negligible.

To demonstrate this error model, the Kuramoto-Sivashinsky (K-S) equation is used. The K-S equation exhibits spatio temporal chaotic behavior for just one spatial dimension, making it an inexpensive surrogate for turbulence. This report considers the K-S equations with the same numerical discretization, parameters, and boundary conditions as in the report by Blonigan & Wang (2014). As for Blonigan & Wang (2014), the instantaneous objective function is the spatial average of the primal solution $u(x,t)$

$$J(u) = \int_0^L u(x)\ dx.$$

Sensitivities were taken with respect to the linear convection parameter c defined by Blonigan & Wang (2014).

Figure 1 shows the effect that ϵ and β have on the sensitivities computed using NILSS for a given solution $u(x,t)$. Noise on the order of 1% of $\|v(t)\|_2$ created by setting $\epsilon = 0.01$ has little effect on the accuracy of NILSS; the sensitivities are distributed tightly around the "true" NILSS sensitivity computed for $\epsilon, \beta = 0.0$. Setting $\epsilon = 0.1$ has some effect, but the tails of the sensitivity distribution remain within $\pm 10\%$ of the true NILSS sensitivity. Noise on the order of $\|v(t)\|_2$ causes sensitivities to be widely distributed. Figure 1(b) shows that the bias β has a much larger impact on the accuracy of the sensitivity.

Figure 2 shows the adjoint magnitude versus time for the same realizations considered for Figure 1. Note that the adjoint remains bounded in all cases, although solutions for higher values of ϵ and β can have some exponential growth. These results and the sensitivity distributions in Figure 1 suggest that NILSS is robust to zero-mean errors on the order of 1–10%. These results also suggest that NILSS can compute accurate

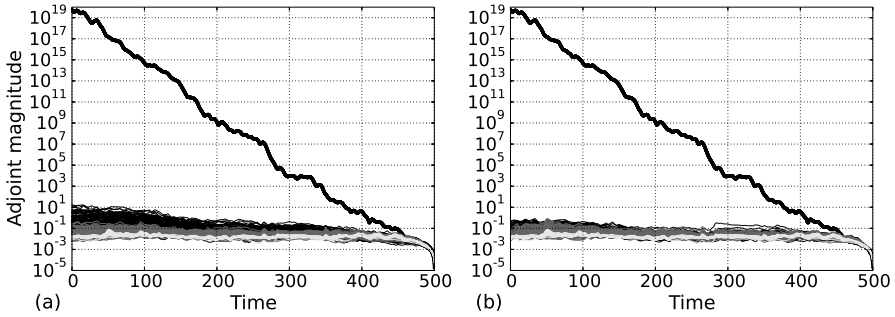

FIGURE 2. NILSS adjoint L2 norms versus time for the adjoints used to generate Figure 1. Thick lines show the magnitude of the conventional adjoint, which grows exponentially backward in time.

sensitivities even if the tangent and adjoint equations are not discretely consistent since the noise term in Eq. (4.1) is not included in the adjoint equations in this study.

Finally, note that increasing the segment size for a fixed value of η will increase the variance of the gradient and cause the adjoint to grow exponentially in magnitude. This means that if the accuracy of the tangent equation is relaxed for NILSS care needs to be taken when selecting a segment size. If it is too long, the error will dominate the tangent solution, and NILSS will fail.

4.2. *Evolution and structure of the unstable subspace*

A necessary precondition for reduced-order models of the tangent equation is that the tangent solution must inhabit a low-dimensional manifold. For NILSS, this means that reduced-order models are of use if the unstable subspace or subsets of the unstable subspace can be approximated on a low-dimensional vector space of some sort. To explore the applicability of projection-based reduced-order models to unsteady tangent equations, this report investigates how the unstable subspace evolves over time.

We use covariant Lyapunov vectors (CLVs) as they correspond to specific time scales (Lyapunov exponent reciprocals) and are norm invariant. The CLVs $\psi^1(u), \psi^2(u), ..., \psi^n(u)$ correspond to each Lyapunov exponent Λ^i and satisfy the evolution equation (Ginelli *et al.* 2007)

$$\frac{d}{dt}\psi^i[u(t)] = \left.\frac{\partial f}{\partial u}\right|_{u(t)} \psi^i[u(t)] - \Lambda^i \psi^i[u(t)]. \tag{4.2}$$

Note that Eq. (4.2) is a linearization of Eq. (2.1) with an additional term $-\Lambda^i\psi^i(u(t))$. Lyapunov exponents and covariant vectors can be computed for numerical simulations. For this report, Lyapunov exponents and CLVs are computed using the algorithms presented by Benettin *et al.* (1980) and Ginelli *et al.* (2007), respectively.

4.2.1. *Evolution of the unstable subspace*

First, we wish to determine if any part of the unstable subspace at time t_i, Q_i is useful at other times $t_j \neq t_i$. In other words, we wish to know whether the orthonormal basis of the unstable space at time t_i can be used to approximate the tangent space at a nearby time t_j. This can be done for a given CLV $\psi^k[u(t_j)]$ by projecting it onto an unstable subspace Q_i and determining which portions of Q_i the CLV is parallel to. This can be

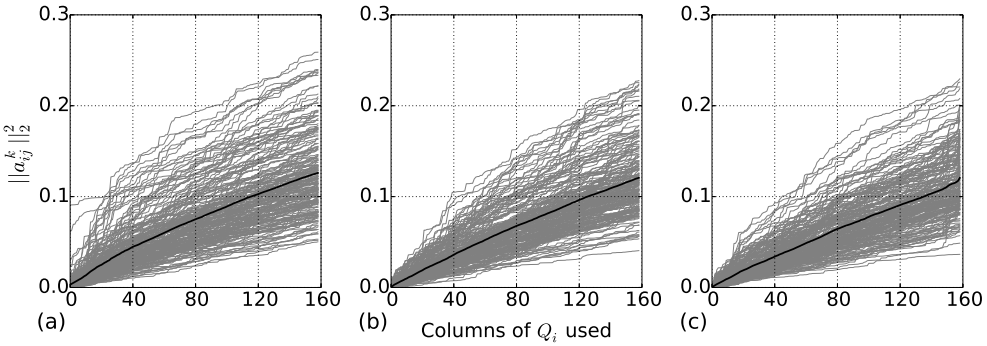

FIGURE 3. Figure showing projection accuracy $[a_{ij}^k]^T a_{ij}^k$ for Λ^1 (a), Λ^{80} (b), and Λ^{159} (c) for $j = i + 1$. The first two modes are unstable ($\Lambda^k > 0$), while the third is nearly neutrally stable ($\Lambda^k \approx 0$). $[a_{ij}^k]^T a_{ij}^k$ for each individual segment are in grey, while their mean is in black.

done by computing the length-p weight vector a_{ij}^k defined as

$$a_{ij}^k = Q_i^T \frac{\psi^k(u(t_j))}{\|\psi^k(u(t_j))\|_2}. \tag{4.3}$$

Since $\psi^k(u(t_j))$ is normalized, the nth component of a_{ij}^k is the cosine of the angle between $\psi^k[u(t_j)]$ and the nth column of Q_i. Therefore, $[a_{ij}^k]^T a_{ij}^k$ indicates how accurately Q_i resolves $\psi^k[u(t_j)]$. If $[a_{ij}^k]^T a_{ij}^k = 1$, then $\psi^k[u(t_j)]$ can be resolved exactly by the span of Q_i. The same idea applies to subsets of columns of Q_i.

To study the unstable subspace for a tubulent flow, we consider the minimal turbulent flow unit for near-wall turbulence originally presented by Jimenez & Moin (1991). Specifically, a $\pi\delta \times 2\delta \times 1.34\pi\delta$ minimal channel with $Re_\delta = 3000$ and $Re_\tau = 140.0$ is considered. Details on the numerical solver and turbulent statistics used are presented by Blonigan (2017).

The CLVs are computed from 200 time segments of length $2.5t_e$, where t_e is the eddy turnover time defined as

$$t_e = \delta / U.$$

Figure 3 shows $[a_{ij}^k]^T a_{ij}^k$ for several CLVs. Overall, it seems that the unstable subspace from one time is not very useful for other times. Even when every column of Q_i is used, $\psi^k(u(t_{i+1}))$ is poorly resolved, since $[a_{ij}^k]^T a_{ij}^k \approx 0.12$ averaged over all time segments. Note that similar results were obtained for the K-S equation. We conclude that knowledge of the unstable subspace at one time does not help with resolving instabilities at other times.

This shows that a projection-based model would not gain much accuracy by using unstable directions computed at other times, even for neutrally stable CLVs, which have relatively slow time scales by definition. The flow instabilities corresponding to various eddies appear, convect, and eventually dissipate along with the eddies themselves. Although there are many eddies of similar size with similar instability mechanisms, they will occur at different locations and times. A linear basis vector will only know one instability on one eddy at one spatial location.

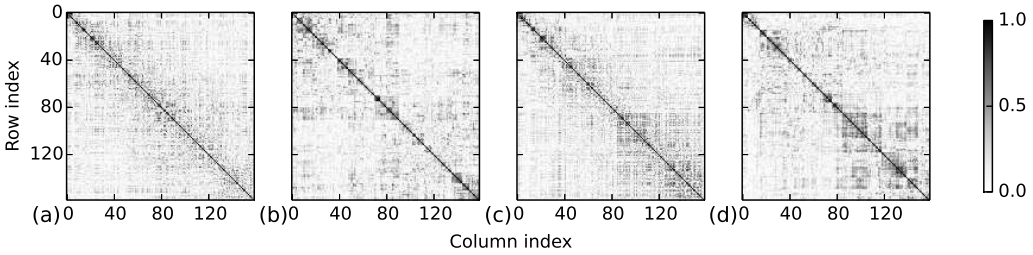

FIGURE 4. Matrix of angle cosines between CLVs for the minimal flow unit at two different times, $t = 100t_e$ (a,b) and $t = 175t_e$ (c,d). Matrices (b) and (d) are reordered using hierarchical clustering.

4.2.2. *Structure of the unstable subspace*

CLVs can also be used to study the structure of the unstable subspace and the strange attractor as a whole. Potential partitioning of the unstable subspace is studied by clustering the angle cosines between unstable CLVs. Clustering groups CLVs with small angles between them together. The small angles indicate some overlap or coupling between the instabilities. A hierarchical clustering approach is used with a Euclidean distance metric and average linkage clustering, also known as the UPGMA algorithm (Sokal & Michener 1958)

Figure 4 shows several snapshots of the cosine matrix for the minimal turbulent flow unit. The clustering algorithm reordered the indices, grouping CLVs with a small angle close to one another. Each snapshot has a number of pronounced groupings, indicated by the darker blocks along the main diagonal. These blocks also typically have very small off-diagonal components, which means that although the CLVs in the group are only separated by small angles, they are nearly orthogonal to the rest of the CLVs. However, it appears that these groups vary in time, so CLVs that are close at one time can be nearly orthogonal at another time.

5. Conclusions

Although chaotic sensitivity analysis approaches such as NILSS are computationally expensive, there might be some ways to reduce their cost. Firstly, the unstable subspace can still be resolved with sufficient accuracy for NILSS with approximate tangent equation solutions. This has been demonstrated by modeling errors in the tangent equation with stochastic forcing when computing the unstable subspace for NILSS. Therefore, coarse discretizations or reduced-order models could be used in place of the full tangent equation in the most expensive part of the NILSS approach, which could reduce the cost of NILSS considerably.

Second, we found that using a fixed linear basis for reduced-order models of the tangent will not work well. For the minimal flow unit, the entire unstable subspace varies significantly with time, and snapshots of it do not form a basis from which the tangent can be resolved accurately. Also, although the CLVs can be grouped together into clusters with small angles between them, these clusters vary significantly over time. Overall, these results suggest that linear dimension reduction is insufficient for the tangent of turbulent flow fields. What is likely needed is nonlinear dimension reduction with some sort of manifold learning approach, which has been used to great effect in computer vision.

Acknowledgments

Resources supporting this work were provided by the NASA High-End Computing (HEC) Program through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center. The authors acknowledge helpful discussions with a number of other summer program participants. We would also like to thank Qiqi Wang for his suggestions on analyzing CLVs and their angles.

Appendix A. Adjoint NILSS algorithms

Algorithm 1 Adjoint NILSS

Inputs Initial condition for the governing equations u_0, Spin-up time t_0, Specified time horizon and checkpoints $t_0, t_1, ..., t_K$

Outputs Sensitivities $d\bar{J}/ds$

Compute $u(t)$ by solving Eq. (2.1) until $t = t_0$.
Set $W_1(t_0) = Q_0$, where Q_0 is some $n \times p$ unitary matrix.
for $i = 1$ to K **do**
    Solve for $u(t)$ from t_{i-1} to t_i using Eq. (2.1).
    Solve for each column of $W_i'(t)$ from t_{i-1} to t_i using Eq. (3.4) with $\eta^p = 0$.
    Compute and save

$$g_{i,1}^T = \frac{1}{\Delta T} \int_{t_{i-1}}^{t_i} \left. \frac{\partial J}{\partial u} \right|_t^T W_i'(t)\, dt.$$

    Compute and save $F_{t_i} W_i'(t_i)$, where $F_{t_i} = f(u(t_i); s) f(u(t_i); s)^T / \|f(u(t_i); s)\|_2^2$.
    Compute $W_i(t_i) = P_{t_i} W_i'(t_i)$, where $P_{t_i} = (I - F_{t_i})$.
    Compute $Q_i R_i = W_i(t_i)$ with a QR decomposition. Save Q_i and R_i.
    Set $W_{i+1}'(t_i) = Q_i$.
end for
for $i = 1$ to K **do**
    Compute and save $g_i^T = g_{i,1}^T + (1/\Delta T) \left(\bar{J} - J(u(t_i); s)\right) F_{t_i} W_i'(t_i)$.
end for
Form and solve the adjoint KKT system Eq. (3.5).
for $i = K$ to 1 **do**
    if $i = K$ **then**
        $\hat{w}(t_K) = x_K - P_{t_K} Q_K \psi_K$,
    else
        $\hat{w}(t_i^-) = P_{t_i}(I - Q_i Q_i^T)\hat{w}(t_i^+) + x_i - P_{t_i} Q_i \psi_i$
    end if
    Solve Eq. (3.6) with terminal condition $\hat{w}(t_i)$ from t_i to t_{i-1}.
end for
Compute sensitivities with the conventional adjoint sensitivity equation, Eq. (3.7).

Note that

$$x_i = \frac{1}{\Delta T} \left(\bar{J} - J(u(t_i); s)\right) \frac{f(u(t_i); s)}{\|f(u(t_i); s)\|_2^2}.$$

REFERENCES

Benettin, G., Galgani, L., Giorgilli, A. & Strelcyn, J. 1980 Lyapunov characteristic exponents for smooth dynamical systems and for Hamiltonian systems; a method for computing all of them. Part 2: Numerical application. *Meccanica* **15**, 21–30.

Blonigan, P. 2017 Adjoint sensitivity analysis of chaotic dynamical systems with non-intrusive least squares shadowing. *J. Comput. Phys.* **348**, 803–826.

Blonigan, P., Fernandez, P., Murman, S., Wang, Q., Rigas, G. & Magri, L. 2016 Toward a chaotic adjoint for LES. *Proceedings of the 2016 Summer Program*, Center for Turbulence Research, Stanford, pp. 385–394.

Blonigan, P. & Wang, Q. 2014 Least squares shadowing sensitivity analysis of a modified Kuramoto–Sivashinsky equation. *Chaos Soliton. Fract.* **64**, 16–25.

Blonigan, P., Wang, Q., Nielsen, E. & Diskin, B. 2018 Least-squares shadowing sensitivity analysis of chaotic flow around a two-dimensional airfoil. *AIAA J.* **56**, 658–672.

Ginelli, F., Poggi, P., Turchi, A., Chaté, H., Livi, R. & Politi, A. 2007 Characterizing dynamics with covariant Lyapunov vectors. *Phys. Rev. Lett.* **99:130601**.

Jimenez, J. & Moin, P. 1991 The minimal flow unit in near-wall turbulence. *J Fluid Mech.* **225**, 213–240.

Keefe, L., Moin, P. & Kim, J. 1992 The dimension of attractors underlying periodic turbulent poiseuille flow. *J Fluid Mech.* **242**, 1–29.

Lea, D., Allen, M. & Haine, T. 2000 Sensitivity analysis of the climate of a chaotic system. *Tellus* **52A**, 523–532.

Leonard, T., Gicquel, L., Gourdain, N. & Duchaine, F. 2015 Steady/unsteady Reynolds-averaged Navier-Stokes and large eddy simulations of a turbine blade at high subsonic outlet mach number. *J. Turbomach.* **137**, 041001.

Ni, A. & Wang, Q. 2017 Sensitivity analysis on chaotic dynamical systems by non-intrusive least square shadowing (NILSS). *J. Comput. Phys.* **347**, 56–77.

Oliver, T., Malaya, N., Ulerich, R. & Moser, R. 2014 Estimating uncertainties in statistics computed from direct numerical simulation. *Phys. Fluids* **26**.

Pulliam, T. & Vastano, J. 1993 Transition to chaos in an open unforced 2D flow. *J. Comput. Phys.* **105**, 133–149.

Sirovich, L. & Deane, A. 1991 A computational study of Rayleigh-Benard convection. Part 2. Dimension considerations. *J. Fluid Mech./* **222**, 251–266.

Sokal, R. & Michener, C. 1958 A statistical method for evaluating systematic relationships. *University of Kansas Science Bulletin* **38**, 1409–1438.

Wang, Q., Hui, R. & Blonigan, P. 2014 Least squares shadowing sensitivity analysis of chaotic limit cycle oscillations. *J Comput. Phys.* **267**, 210–224.