Newsjournal of the Society for Industrial and Applied Mathematics sinews.siam.org

siam news

Volume 55/ Issue 5
June 2022

Dynamics-based Machine Learning for Nonlinearizable Phenomena

Data-driven Reduced Models on Spectral Submanifolds

By George Haller, Shobhit Jain, and Mattia Cenedese

Machine learning (ML) has been an inspiring development for all areas of applied science, with numerous success stories in static learning environments like image, pattern, and speech recognition. Yet effective modeling of dynamical phenomena—such as nonlinear vibrations of solids and transitions in fluids—remains a challenge for ML, which tends to produce overly complex and uninterpretable dynamic models that are not reliable outside of their training range. A recent approach, however, integrates advanced dynamical systems concepts into elementary ML, ultimately yielding fast and accurate reduced-order models for nonlinear dynamics.

The idea—which we call *dynamics-based machine learning* (DBML)—is to learn models directly from phase space structures that are inferred from data. Systems with very different physics often display the same key invariant sets in their phase spaces; instead of fitting models to individual trajectories (which are sensitive to perturbations and parameter changes anyway), robust reduced-order modeling should therefore target structurally stable invariant sets. DBML focuses specifically on identifying the dynamics of ubiquitous, low-dimensional attracting invariant manifolds, which were first noted in the nonlinear vibrations literature [6]. Subsequent work in dynamical systems theory independently established the existence and properties of these manifolds, even for infinite-dimensional systems [1]. The forthcoming formulation, a higher-dimensional computational algorithm, and a data-driven implementation of these results have only appeared very recently [2-4].

DBML assumes the existence of at least one stationary state $\mathcal{M}_0$ for a dynamical system, which we take here to be finite dimensional for simplicity. To further simplify the situation, we only consider the case wherein $\mathcal{M}_0$ is an attracting fixed point; similar results hold for repelling fixed points, periodic orbits, and quasiperiodic steady states. The linearized dynamical system at $\mathcal{M}_0$ will admit eigenspaces E_j that are spanned by generalized eigenvectors of its jth distinct eigenvalue λ_j. We can order these eigenspaces by their increasing real parts, so that $\mathrm{Re}\lambda_j < \mathrm{Re}\lambda_{j+1}$. As a consequence, solutions of the linearized system within E_j decay to the fixed point $\mathcal{M}_0$ increasingly quickly as the index j grows.

By grouping some of the E_j eigenspaces together if necessary, we can build a hierarchy $E^1 \subset E^2 \subset \ldots$ of *spectral subspaces* (see Figure 1a). We construct these spectral subspaces from eigenspaces in a manner that ensures that all E^j are non-resonant. In particular, no positive-integer linear combination of the eigenvalues in E^j should equal any eigenvalue that falls outside of E^j.

See **Machine Learning** *on page 4*

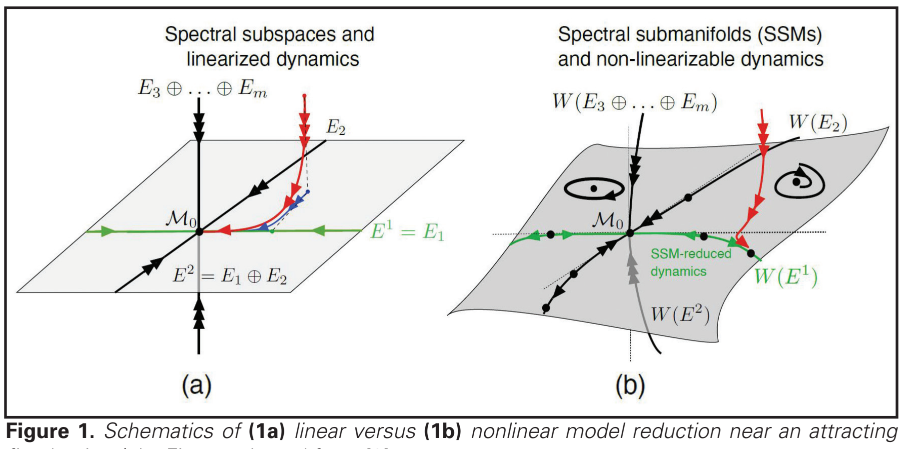

Figure 1. *Schematics of* **(1a)** *linear versus* **(1b)** *nonlinear model reduction near an attracting fixed point* $\mathcal{M}_0$*. Figure adapted from [2].*

How to Boost Your Creativity

By Nicholas J. Higham and Dennis Sherwood

As a *SIAM News* reader, you are a creative person — you would not have gotten where you are today without being creative. However, you may not understand exactly how your creativity works or how to "turn it on" when you need it. You also might not know how to train others to be creative.

In this article, we describe the basic idea behind a six-step process for creativity that Dennis Sherwood developed over the last 20 years. We have used this process together in creativity workshops during the last decade, working with groups ranging from the numerical analysis group at the University of Manchester to the SIAM leadership.

Our approach is based on the insight that creativity is not so much about creating something totally new as about identifying something *different*. The search for something different is much easier than the search for something new, for "different" means "different from now" and "now" is visible all around us. So if we observe "now" very carefully, we might notice some feature of "now" that might be different and ideally better too, and from this feature an idea might spring. Creativity therefore does not necessarily require an act of genius, or a lightning strike out of the blue. Rather, good ideas can be discovered as the result of detailed observation coupled with curiosity, and can follow a systematic process that can be applied in any circumstance.

Two of the key steps in our procedure are to write down every feature of the focus of attention (which could be a mathematical problem or something else entirely), and then ask "How might this be different?" for each one. Here we provide a glimpse into the process with an old and familiar example: iterative refinement for improving an approximate solution to a linear system $Ax = b$, where A is a square, nonsingular matrix. The basic algorithm in its original form is as follows, and we assume that it is carried out in double-precision floating-point arithmetic:

1. Solve $Ax = b$.
2. Compute $r = b - Ax$ in quadruple precision.
3. Solve $Ad = r$.
4. Update $x \leftarrow x + d$.

Repeat from step 2 if necessary.

Participant at a creativity workshop adds ideas to a flip chart. Photo courtesy of Dennis Sherwood.

James H. Wilkinson programmed iterative refinement in 1948 using LU factorization with partial pivoting for the solves in steps 1 and 3. For step 2, he took advantage of the ability of his computer—the Pilot Automatic Computing Engine at the National Physical Laboratory—to accumulate inner products in quadruple precision at no extra cost.

Iterative refinement became popular and was implemented in this way for the next 25 years or so. Several textbooks from the 1960s and 1970s made statements such as "It is absolutely essential that the residuals be computed in extra precision," and the method seemed to be set in stone. However, every aspect of iterative refinement is amenable to the question "How might this be different?", and the answers to this question have yielded a panoply of different versions of the method.

Here is a thumbnail sketch of some iterative refinement variants, each of which is identified by the feature that distinguishes it from the aforementioned version. Specific references for these and other developments are given in [1, 2].

Precision of the Residual

The residual r does not need to be computed with extra precision, at least not if the aim is to improve the backward stability rather than the accuracy. This was realized in the 1970s, by which time most computers could no longer accumulate inner products in quadruple precision for free. This finding opened up the possibility of using a somewhat unstable solver.

Precision of the Factorization

The algorithm still works if step 1 uses an LU factorization that is computed in single precision, as long as A is not nearly singular to single precision. This observation was made in the 2000s and was important because processors were appearing on which single precision arithmetic was much faster than double precision arithmetic.

See **Creativity** *on page 3*

Nonprofit Org
U.S. Postage
PAID
Permit No 360
Bellmawr, NJ

siam.
SOCIETY for INDUSTRIAL and APPLIED MATHEMATICS
3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 USA

Machine Learning

Continued from page 1

Each E^j thus contains linearized solutions that do not exchange energy via resonances with higher members of the spectral subspace hierarchy.

The subspace E^j serves as an observed attractor for typical linearized trajectories until the components of those trajectories in $E^j - E^{j-1}$ die out. At that point, E^{j-1} becomes the observed attractor (see Figure 1a, on page 1). The reduced dynamics on E^j therefore provide the best possible reduced model of the linearized dynamics if we wish to filter out transients that are associated with all stronger decay exponents $\mathrm{Re}\,\lambda_\ell$ for $\ell > j$.

The fundamental result of *spectral submanifold* (SSM) theory is that this hierarchy of observed linear attractors also persists in a smoothly deformed form within the full nonlinear dynamical system. Specifically, a nested family of SSMs $W(E^1) \subset W(E^2) \subset \ldots$ exists such that $W(E^j)$ is invariant under the full dynamics, has the same dimension as E^j, and is tangent to E^j at the steady state $\mathcal{M}_0$. These SSMs are not unique; they share their invariance, dimensionality, and tangency to E^j with infinitely many other manifolds. Under the addition of small periodic or quasiperiodic forcing, both $\mathcal{M}_0$ and its SSMs persist smoothly and inherit the time dependence of the forcing.

Therefore, SSM-reduced dynamics provide a hierarchy of mathematically exact low-dimensional models for nonlinearizable behavior — even with the addition of moderate external forcing. Such behavior includes coexisting steady states, transitions among them, and chaotic dynamics. SSM-reduced models can be computed in seconds or minutes and reveal the details of nonlinearizable, damped-forced responses in mechanical systems with tens or even hundreds of thousands of degrees of freedom. For example, the red curve in Figure 2 traces an accurate and highly accelerated prediction of forced response from a two-dimensional reduced model on $W(E^1)$ for a 267,840-dimensional finite element model of an aircraft wing. Such a numerical prediction is currently impossible for even the most advanced numerical continuation packages [4].

The SSMTool[1] computations in Figure 2 require explicit knowledge of nonlinearities in the governing equations, which is not available from commercial finite element codes. Even the evaluation of functions that implicitly define the nonlinearities is costly. The prohibitive expense for long-term simulations of individual trajectories means that model reduction is unavoidable.

One possible workaround is a fully data-driven algorithm for SSM construction [2]. This algorithm—which is implemented in an open-source MATLAB package called SSMLearn[2]—uses data to identify the dimension and spectrum of the dominant spectral subspace E^j. The procedure then utilizes regression to reconstruct the SSM in the observable space and computes a sparse normal form for the SSM-reduced dynamics.

This approach yields previously unthinkable computational speed-ups for dynamic finite element simulations. One can simply learn the unforced normal form on SSMs from a small number of decaying, unforced trajectories, then use these low-dimensional models to predict full bifurcation curves of the forced response without any simulation. The blue curve in Figure 2 is an example of this type of nonintrusive, data-driven model reduction, which yields remarkably close agreement with the exact analytic predictions from SSMTool. Here, SSMLearn was trained on a single unforced trajectory and predicted the full forced response curve in only five minutes.

SSM-based model reduction has multiple other uses as well. It is equally applicable to experimental data with arbitrary physics, such as sloshing dynamics in surface-wave experiments that are relevant in the design of tanks on cargo ships and commercial trucks (see Figure 3). Data-driven SSM reduction also provides low-dimensional reduced models for global transitions in infinite-dimensional nonlinear evolution equations. Such transitions occur, for example, in the Navier-Stokes equations for planar Couette flows, which admit multiple steady states beyond their stable, constant-shear base state (see Figure 4).

Here we have illustrated a physically diverse group of dynamical data sets from which DBML constructs accurate and predictive reduced-order models for nonlinearizable dynamics on SSMs. These dynamics display coexisting stable and unstable steady states with transitions among them, which cannot be simultaneously captured by a linear model. Promising ongoing extensions of SSM theory to more general $\mathcal{M}_0$ sets, external forcing types, and non-smooth effects will further enhance the power of DBML.

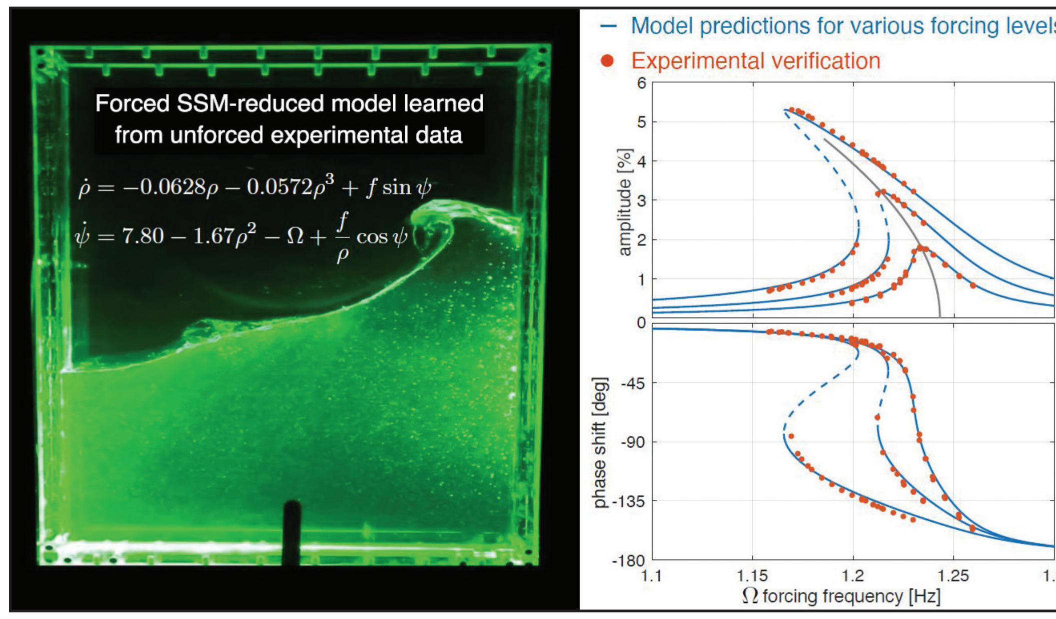

Figure 3. *Predictions for the forced horizontal oscillations of a sloshing fluid's center of mass by a model that is trained on unforced experimental data. Figure adapted from [2] and left image courtesy of Kerstin Avilla.*

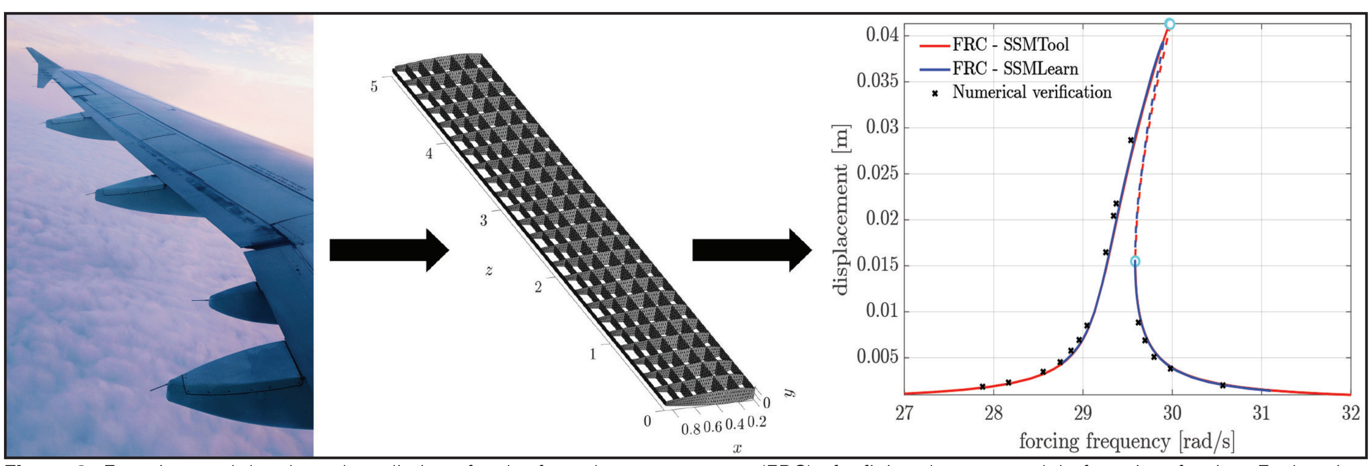

Figure 2. *Equation- and data-based predictions for the forced response curve (FRC) of a finite element model of an aircraft wing. Each trajectory integration in the forced finite element model (black cross) requires approximately four days to cover roughly 20 seconds of physical model time. In contrast, an equation-driven, spectral submanifold (SSM)-based prediction by SSMTool [4] (red curve) for the full FRC takes about 40 minutes. Finally, a purely data-driven FRC prediction by SSMLearn [2] (blue curve) that is trained on a single unforced trajectory takes about five minutes. Dashed portions of the SSM-based FRC predictions, which indicate unstable periodic response, are unavailable to direct numerical integration. Figure courtesy of the authors.*

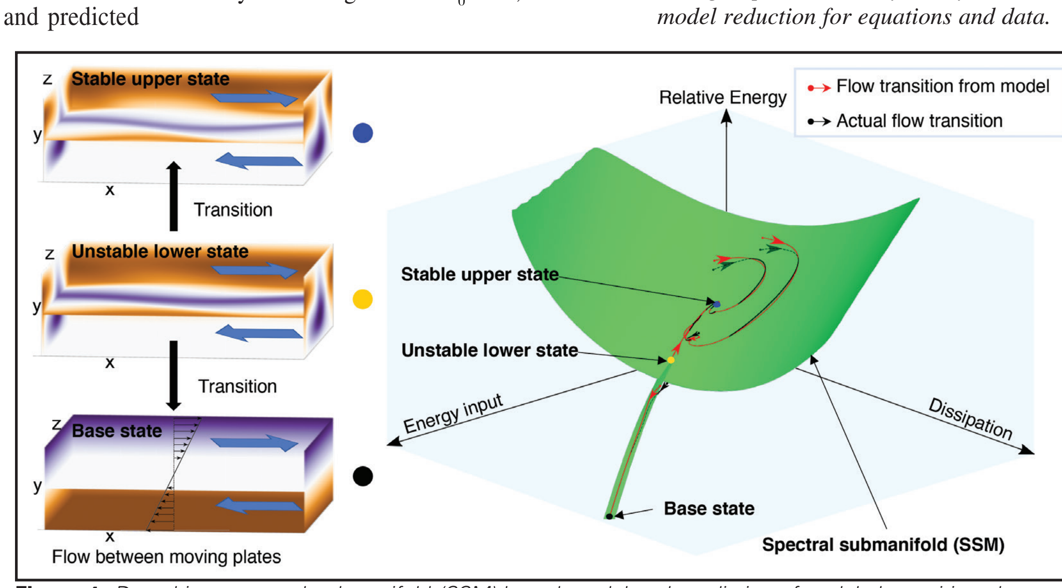

Figure 4. *Data-driven spectral submanifold (SSM)-based model and predictions for global transitions between stable and unstable states in a Couette flow. Figure adapted from [5].*

References

[1] Cabré, X., Fontich, E., & de la Llave, R. (2003). The parameterization method for invariant manifolds I: Manifolds associated to non-resonant subspaces. *Indiana Univ. Math. J.*, *52*(2), 283-328.

[2] Cenedese, M., Axås, J., Bäuerlein, B., Avila, K., & Haller, G. (2022). Data-driven modeling and prediction of non-linearizable dynamics via spectral submanifolds. *Nat. Commun.*, *13*, 872.

[3] Haller, G., & Ponsioen, S. (2016). Nonlinear normal modes and spectral submanifolds: Existence, uniqueness and use in model reduction. *Nonlin. Dynam.*, *86*, 1493-1534.

[4] Jain, S., & Haller, G. (2021). How to compute invariant manifolds and their reduced dynamics in high-dimensional finite-element models. *Nonlin. Dynam.*, *107*, 1417-1450.

[5] Kaszás, B., Cenedese, M., & Haller, G. (2022). Dynamics-based machine learning of transitions in Couette flow. Preprint, *arXiv:2203.13098*.

[6] Shaw, S.W., & Pierre, C. (1993). Normal modes for non-linear vibratory systems. *J. Sound Vibr.*, *164*(1), 85-124.

George Haller is a professor of mechanical engineering at ETH Zürich, where he holds the Chair in Nonlinear Dynamics. His group works on various aspects of nonlinear dynamical systems that are defined by data sets. Shobhit Jain and Mattia Cenedese are postdoctoral researchers in this group, where they study nonlinear model reduction for equations and data.

[1] https://github.com/haller-group/SSMtool

[2] https://github.com/haller-group/SSMLearn

Truth to Power

Continued from page 2

as money in media circles. Speaking truth to power can thus have consequences, like the loss of funding or access to necessary information. Either of these outcomes can damage a career — as well as one's ability to make a living.

"Back Off, Man, I'm a Scientist"

"It's difficult when those in power are so ideologically removed from one's own principles or worldview," Al-Khalili said. "The model of just saying, 'We are the experts, you are the empty vessels to be filled with our knowledge and wisdom,' is not going to work."

Instead of the *Ghostbusters* "Back off, man, I'm a scientist!" approach, Al-Khalili argued that scientists must explain their methodologies and expose their own humanity. Krummel expanded upon this viewpoint. "[Media] has popularized the scientist as white coat-wearing, bespectacled, and very precise," he said. "And there are realities to that. But many of us love the outdoors, many of us go to church and are spiritual — all of this stuff is not part of the popular perception."

Such commonalities can help bridge the divide. Krummel added that scientists should speak from the heart, verbalize how they feel about the issues at hand, and acknowledge that audiences respond to who they are as people, not just to their facts.

In short, science can provide the necessary evidence to oppose ideologically-driven assaults on climate change, LGBTQ+ rights and healthcare, the U.S. Supreme Court's recent leaked stance on abortion, and other important issues. At the same time, the shared humanity between scientists and non-scientists may be a more effective way for science to speak truth to power.

References

[1] Oreskes, N., & Conway, E.M. (2010). *Merchants of doubt: How a handful of scientists obscured the truth on issues from tobacco smoke to global warming*. New York, NY: Bloomsbury.

Matthew R. Francis is a physicist, science writer, public speaker, educator, and frequent wearer of jaunty hats. His website is BowlerHatScience.org.