LEARNING LOW-DIMENSIONAL FEATURE DYNAMICS OF TURBULENT FLOWS USING DEEP CONVOLUTIONAL RECURRENT AUTOENCODERS

Allocation: Exploratory/50 Knh
PI: Maciej Balajewicz[1]
Collaborator: Francisco Gonzalez[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Dynamical systems are used to describe the rich and complex evolution of many real-world processes. Modeling the dynamics of physical, engineering, and biological systems is thus of great importance in their analysis, design, and control. For systems with existing models based on first principles, high-fidelity solutions are possible through direct numerical simulations. However, these generally yield sets of equations with approximately 10^{6-9} degrees of freedom. Even with recent advances in computational power, solving these high-fidelity models is computationally intractable for multiquery and time-critical applications such as design optimization, uncertainty quantification, and model predictive control. Moreover, some systems may have an abundance of data but could lack the governing laws necessary for accurate modeling. Motivated by this problem, we seek to develop deep learning-based model-reduction approaches, wherein both the identification and evolution of low-dimensional features are learned from numerical and experimental data sets.

RESEARCH CHALLENGE

In recent years, the rise of machine learning and big data have driven a shift in the way complex spatiotemporal systems are modeled. The abundance of data has facilitated the construction of so-called data-driven models of systems lacking high-fidelity governing laws. In areas where high-fidelity models do exist, data-driven methods have become an increasingly popular approach to tackle previously challenging problems wherein solutions are learned from physical or numerical data [1].

In model reduction, machine learning strategies have recently been applied to many remaining challenges, including learning stabilizing closure terms in unstable proper orthogonal decomposition (POD)–Galerkin models and data-driven model identification for truncated generalized POD coordinates [2,3]. A more recent approach considers learning a set of observable functions spanning a Koopman invariant subspace from which low-order linear dynamics of nonlinear systems are modeled [4]. While many of these approaches show great promise, a number of significant issues remain. Notably, the issue of scalability arises when considering training deep neural networks on large-scale simulation data.

To avoid this curse of dimensionality, we instead propose a deep learning method that combines important innovations in dimensionality reduction and arbitrary dynamics modeling to perform robust deep learning-based model reduction. First, a deep neural network architecture called a convolutional autoencoder (Fig. 1a) is used to learn low-dimensional, abstract features of the high-dimensional input data. A modified version of a long short-term memory (LSTM) recurrent neural network (Fig. 1b) is then used to learn the *a priori* unknown dynamics of these features. Both networks are trained jointly in an end-to-end fashion, resulting in a completely data-driven model that offers significant advantages over both linear model reduction strategies and vanilla implementations of neural-network-based model-reduction approaches.

METHODS & CODES

The model proposed in this work consists of a four-layer strided convolutional encoder network followed by a two-layer dense encoder network, both of which learn at each layer a lower-dimensional abstract representation of the input data. This results in a low-dimensional feature vector, $h \in \mathbb{R}^k$, which can be thought of as a nonlinear, location-invariant generalization of the generalized POD coordinates. To efficiently evolve these features, the LSTM network is designed to scale with the reduced dimension, relying only on the current state to make a future state prediction. Both the LSTM and encoder networks use the decoder network, which consists of a two-layer dense network followed by a four-layer strided convolutional transposed network, to reconstruct the full state from the low-dimensional feature vector.

The code used in this project is written in Python using TensorFlow, Google's open source library for building, training, and serving deep neural network models, which utilizes the CUDA CuDNN deep learning library for acceleration with NVIDIA GPUs [5]. The model is trained using ADAM, a variant of stochastic gradient descent.

RESULTS & IMPACT

We trained our deep convolutional recurrent autoencoder model on a number of illustrative examples. Here, we restrict our attention to the problem of a statistically stationary lid-driven cavity flow at a high Reynolds number. In particular, the Navier–Stokes equations in streamfunction–vorticity formulation are discretized in space using Chebyshev polynomials and integrated in time using a semi-implicit second order scheme at a Reynolds number of $Re=2.5\times10^4$. At such high Reynolds numbers, this flow exhibits complex spatiotemporal behavior and serves as a well-known benchmark for validation of reduced-order models.

Next, we assembled a data set of finite-time solution snapshot sequences where each snapshot consists of feature-scaled streamfunction fluctuations around the temporal mean. At every step in the training procedure, every individual snapshot is reduced to a low-dimensional feature vector and reconstructed, and the evolution of the feature vector is compared against the current optimal compression. Fig. 2 compares final evolved output feature-scaled streamfunction fluctuations and the corresponding u-velocity and vorticity fields at three different stages during training using a feature vector $h \in \mathbb{R}^k$, with $k=8$. With no training, the model simply outputs noise from the random initialization of the model parameters. At 7,000 training steps, the model begins to learn the evolution of the system, and by 600,000 training steps the model nearly captures the exact solution.

In this work, we have successfully demonstrated the feasibility of using deep neural network architectures for learning and evolving low-dimensional features of high-dimensional systems through the example of a high-Reynolds-number lid-driven cavity flow. The incorporation of machine and deep learning strategies for constructing smarter and more efficient reduced-order models is still a nascent field, but it is one that could have a significant impact on areas ranging from design optimization, uncertainty quantification, and model predictive control. To this end, we are pursuing a number of different directions including incorporating physics-based constraints to learning low-order dynamics, and learning low-order feature dynamics from heterogeneous data sets.

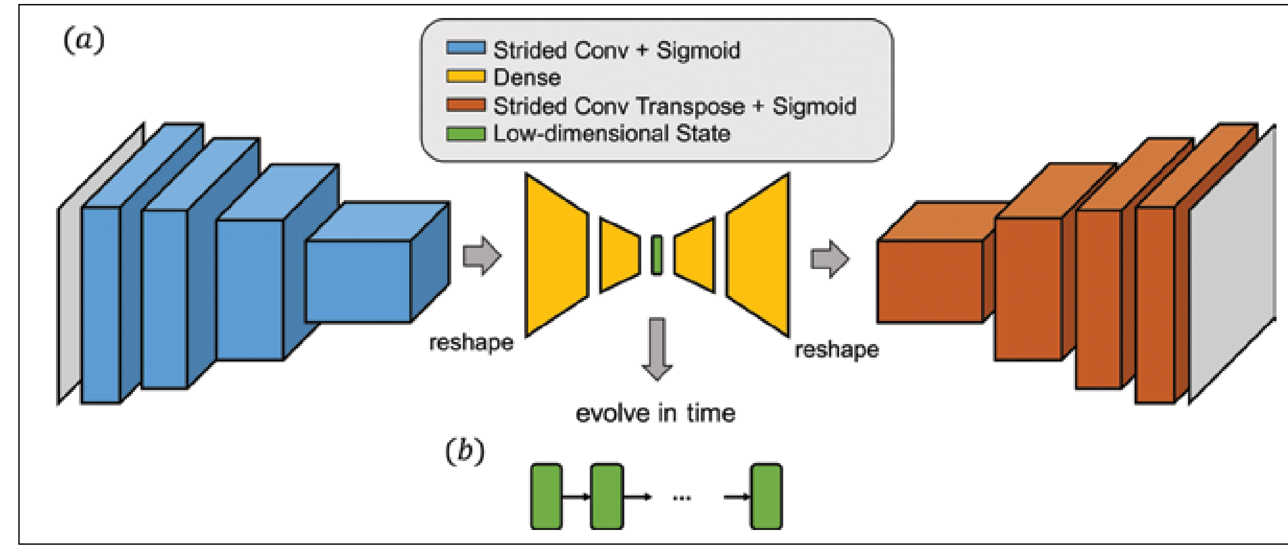

Figure 1: (a) The deep convolutional autoencoder model proposed in this work. In this model, a 2D input snapshot is fed on the left and is processed down to a low-dimensional state (or feature vector). (b) The long short-term memory (LSTM) then evolves this feature vector, where the full state is reconstructed with the decoder portion of the convolutional autoencoder.

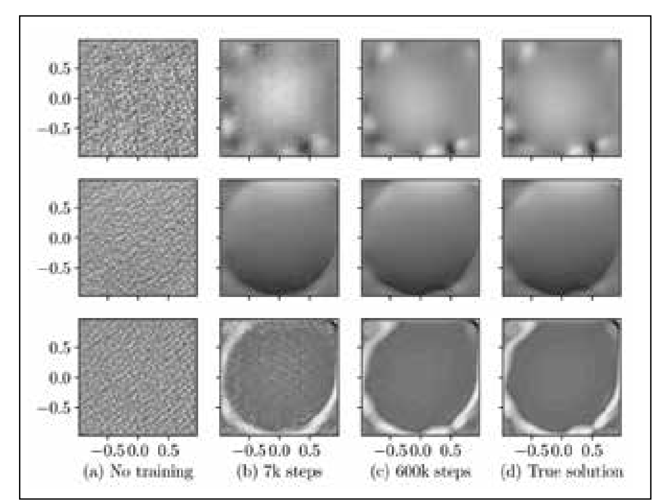

Figure 2: Evolution of the model output feature-scaled streamfunction fluctuations (top row), its corresponding u-velocity field (middle row), and vorticity field (bottom row) at three different stages during training: (a) no training, (b) after 7,000 training steps, and (c) after 600,000 training steps. The exact solutions are shown in (e).

WHY BLUE WATERS

Training deep neural network models is an inherently data-intensive process. With larger simulations and more sophisticated sensors, there is no shortage of data from which deep learning models of physical systems can be trained. The petascale resources available via Blue Waters, and in particular its large number of GPU-equipped nodes and fast shared parallel storage, have made developing and training deep neural network-based reduced-order models possible.