Nonlinear reduced-order modeling: Using machine learning to enable extreme-scale simulation for many-query problems

Kevin Carlberg, Sandia National Laboratories

Physics-based modeling and simulation has become indispensable across many applications in engineering and science, ranging from nuclear-weapons design to monitoring national critical infrastructure. However, as simulation is playing an increasingly important role in scientific discovery, decision making, and design, greater demands are being placed on model fidelity. This high fidelity necessitates modeling fine spatiotemporal resolution, which can lead to extreme-scale, nonlinear dynamical-system models whose simulations consume months on thousands of computing cores. Further, most scientific applications (e.g., uncertainty quantification, design optimization) are many query in nature, as they require the (parameterized) model to be simulated thousands of times. This leads to a computational barrier: the computational cost of high-fidelity simulations renders them impractical for many-query problems.

In this talk, I will present several advances in the field of nonlinear model reduction that exploit simulation data to overcome this barrier. These methods combine concepts from machine learning, computational mechanics, optimization, and goal-oriented adaptivity to produce low-dimensional reduced-order models (ROMs) that are 1) accurate, 2) low cost, 3) structure preserving, 4) reliable, and 5) certified. First, I will describe least-squares Petrov–Galerkin projection, which leverages subspace identification and optimal projection to ensure accuracy. Second, I will describe the sample mesh concept, which employs empirical regression and greedy-optimal sensor-placement techniques to ensure low cost. I will also describe novel methods that exploit time-domain data to further reduce computational costs. In addition, I will describe deep model reduction, which employs convolutional autoencoders from deep learning to substantially reduce the ROM dimensionality further. Third, I will present a technique that ensures the ROM is globally conservative in the case of finite-volume discretizations, thus ensuring structure preservation. Fourth, I will describe ROM h-adaptivity, which employs concepts from adaptive mesh refinement to ensure that the ROM is reliable, i.e., it can satisfy any prescribed error tolerance. Finally, I will present machine-learning error models, which apply regression methods (e.g., feedforward neural networks) from machine learning to construct a stochastic model for the ROM error; this quantifies the ROM-induced epistemic uncertainty and provides a mechanism for certification.