Dynamical Reduced-Order Models for High-Dimensional Systems

by

Aaron Charous

Sc.B., Brown University (2019)

S.M., Massachusetts Institute of Technology (2021)

Submitted to the Department of Mechanical Engineering and
Center for Computational Science & Engineering
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in
Mechanical Engineering and Computational Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2023

Author ....

Department of Mechanical Engineering
Center for Computational Science & Engineering
June 1, 2023

Certified by ....

Pierre F.J. Lermusiaux
Professor, Department of Mechanical Engineering
Thesis Supervisor

Accepted by ....

Nicolas Hadjiconstantinou
Professor, Department of Mechanical Engineering
Chairman, Department Committee on Graduate Theses

Accepted by ....

Youssef M. Marzouk
Professor, Department of Aeronautics and Astronautics
Co-Director, Center for Computational Science & Engineering

Dynamical Reduced-Order Models for High-Dimensional Systems

by

Aaron Charous

Submitted to the Department of Mechanical Engineering
and Center for Computational Science & Engineering
on June 1, 2023, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in
Mechanical Engineering and Computational Science

Abstract

Advances in computational power have brought the possibility of realistically modeling our world with numerical simulations closer than ever. Nevertheless, our appetite for higher fidelity simulations and faster run times grows quickly; we will always grasp at what is just beyond our computational reach. No matter if we seek to understand biological, chemical, or physical systems, the bottleneck of scientific computing is almost always the same: high dimensionality. The goal of reduced-order modeling is to reduce the number of unknowns in a system while minimizing the loss of accuracy in the approximate solution. Ideal model-order reduction techniques are optimal compromises between computational tractability and solution fidelity. While there are plenty of such techniques to choose from, their widespread adoption remains to be seen due to several persistent challenges. Many methods are intrusive and difficult to implement, lack traditional numerical guarantees such as convergence and stability, and cannot adapt to unforeseen dynamics. We seek to promote the adoption of reduced-order models (ROMs) by creating non-intrusive, efficient, and dynamically adaptive algorithms that maintain the essential features and numerical guarantees of their full-order counterparts.

In this thesis, we derive and apply algorithms for *dynamical* reduced-order models. Many model-order reduction approaches project dynamical systems onto a fixed subspace obtained from either a simplification of the original equations, a set of known functions such as orthogonal polynomials, or a reduced basis of full-order simulations computed offline. However, if the true system exits the span of the prescribed subspace, such approaches quickly accumulate large errors. In contrast, dynamical ROMs adapt their subspaces as the system evolves. Geometrically, this amounts to integrating a dynamical system along a nonlinear manifold embedded in a full-order Euclidean space. We develop schemes that not only change subspaces at each discrete time step, but that change the subspace in between time steps for improved accuracy. Even further, our numerical schemes automatically detect when the dy-

namics depart the nonlinear manifold and may jump to a new nonlinear manifold that better captures the system state. For concreteness, we focus on a reduced-order modeling technique called the dynamical low-rank approximation (DLRA), a discrete analogue to the dynamically orthogonal (DO) differential equations. The DLRA evolves a low-rank system in time (or range) as an approximation to a full-rank system, and in contrast to many methods, the DLRA does not require an offline stage where full-order simulations are computed. It is also agnostic to the source of high dimensionality, whether it be the high resolution required, the large domain, or the stochasticity of the problem. These features make it a versatile tool suitable for a wide variety of problems. We evaluate, verify, and apply our new dynamical reduced-order models and schemes to a varied set of dynamical systems, including stochastic fluid flows and waves, videos and their dynamic compression, realistic ocean acoustics and underwater sound propagation with dynamic coordinate transforms, and stochastic reachability and time-optimal path planning.

The majority of this work is devoted to new adaptive integration schemes for the DLRA. We start by introducing perturbative retractions, which map arbitrary-rank matrices back to a manifold of fixed-rank matrices. They asymptotically approximate the truncated singular value decomposition at a greatly reduced cost while guaranteeing convergence to the best low-rank approximation in a fixed number of iterations. From these retractions, we develop the dynamically orthogonal Runge-Kutta (DORK) schemes, which change the subspace onto which the system's dynamics are projected in between time steps. The DORK schemes are improved by using stable, optimal (so) perturbative retractions, resulting in the so-DORK schemes. They are more efficient, accurate, and stable than their predecessors. We also introduce gradient-descent (gd) retractions and the gd-DORK schemes, which tend to converge rapidly to the best low-rank approximation by recursively applying retractions. The DORK schemes may be made rank-adaptive and robust to rank overapproximation with either a pseudoinverse or by changing the gauge of the integration scheme. While the pseudoinverse technique accumulates slightly more error, it preserves mode continuity, a feature that changing the gauge lacks. Next, we derive an alternating-implicit (ai) linear low-rank solver, which is used to create ai-DORK schemes. The ai-DORK schemes are a general-purpose family of implicit integration schemes that have the same algorithmic complexity as explicit schemes (provided some conditions on the dynamics), which vastly broadens the scope of problems that can be solved with the DLRA. This relieves stringent time-step restrictions and enables the DLRA to handle stiff systems. Furthermore, we develop a piecewise polynomial approximation using adaptive clustering in order to handle non-polynomial nonlinearities in reduced-order models. We thoroughly test these numerical schemes on well-conditioned and ill-conditioned matrix differential equations; data-driven dynamical systems including videos; Schrödinger's equation; a stochastic, viscous Burgers' equation; a deterministic, two-dimensional, viscous Burgers' equation; an advection-diffusion partial differential equation (PDE); a nonlinear, stochastic Fisher-KPP PDE; nonlinear, stochastic ray tracing; and a nonlinear, stochastic Hamilton-Jacobi-Bellman PDE for time-optimal path planning. We find that the reduced-order solutions may be made arbitrarily accurate using rank-adaptive dynamical schemes that automatically track the true rank of the full-order simulation, and nonlinearities may be well-approximated by dynamically increasing the number of stochastic clusters, all at a greatly reduced computational cost.

In addition to DORK schemes, we create a tailor-made low-rank integration scheme for the narrow-angle parabolic wave equation called the low-rank split-step Fourier method. Acoustic simulations are often bottlenecked by the Nyquist criterion, which insists that we sample spatially at least twice per wavelength. To address this, our low-rank split-step Fourier method has an algorithmic complexity that scales sublinearly in the number of classical degrees of freedom, enabling vastly larger computational domains and higher frequencies. We demonstrate its efficacy on realistic ocean acoustics problems in Massachusetts Bay with sound speed fields obtained from our high-resolution ocean primitive equations modeling system. In comparing the low-rank and full-rank simulations, we demonstrate that the dynamical low-rank method captures the full-rank features including three-dimensional acoustic energy propagation in complex ocean fields with internal waves and rapidly varying bathymetry.

Lastly, with tools from machine learning, we introduce learnable and automatically differentiable coordinate transforms. The compressibility of a system heavily depends on the choice of coordinates, and frequently a coordinate system is chosen for its simplicity rather than its efficiency. Our novel coordinate transforms are determined in a hands-off manner by minimizing a cost function that includes the environmental data expressed in terms of the non-constant coefficients and initial conditions of a PDE. Not only do we automatically obtain Jacobians and Hessians of the transforms, we also find coordinate systems that reduce the rank of solutions to PDEs. This improves the accuracy of the DLRA for the same cost as a typical low-rank simulation, and it accelerates the convergence in rank to the full-order solution. The coordinate transforms also enable low-rank domain decomposition, which is particularly useful in ocean acoustics where the water-seabed interface is discontinuous. We demonstrate this methodology on a first-order PDE with advection and a second-order PDE, the parabolic wave equation, using two examples. We first show acoustic propagation along a three-dimensional wedge and compare the accuracy of solutions computed in the original and transformed coordinate systems. We then show acoustic propagation in a realistic ocean environment over Stellwagen Bank in Massachusetts Bay with a dynamic coordinate transform.

Thesis Supervisor: Pierre F.J. Lermusiaux
Title: Professor, Department of Mechanical Engineering

Acknowledgments

The first time I met with my advisor, Professor Pierre Lermusiaux, I told him I didn't like fluid mechanics. For the unacquainted, the MSEAS group Pierre leads focuses on understanding phenomena in the ocean. This didn't stop Pierre from offering me a spot in MSEAS, and it didn't stop me from taking it. From the start, Pierre has offered me guidance for navigating the lab, MIT, and the academic world at large. At first, I didn't understand every piece of advice Pierre gave me. I certainly didn't understand every research idea he proposed. But as time went on, I started to appreciate his advice more and more; it became clear that Pierre had my best interests in mind even when I couldn't see the full picture. Pierre helped me form research ideas, but he also gave me room to breathe, room to think on my own, room to be creative, room to pursue what I found interesting. From Pierre, I have learned how to collaborate on a very abstract idea and then map it down to a very concrete implementation. I have learned to see the connections between all the research going on inside MSEAS (from our beloved group meetings) and outside our group. I have learned that academic research is, like most important things in life, about relationships with people, and I am so glad to have had such a meaningful one with Pierre.

I would like to thank Dr. Chris Mirabito and Dr. Patrick J. Haley, Jr., for all of the help over these years. You two are the pillars of MSEAS, not only supporting the group's core research functions, but also adding humor and good spirit. In addition, thank you to my committee, Professor Laurent Demanet, Professor Youssef M. Marzouk, and Professor Anthony T. Patera. With an outside perspective, you each have provided invaluable feedback that has strengthened my thesis and often made me sit down to think closely about what I am doing and, more importantly, why. I am grateful to the Office of Naval Research for partial support under grants N00014-19-1-2693 (IN-BDA) and N00014-19-1-2664 (Task Force Ocean, DEEP-AI). In addition, the administrative support at MIT has been first-class. I would like to thank Lisa Maxwell, Kate Nelson, Leslie Regan, Una Sheehan, and Saana McDaniel for all of the rapid responses and support these past years.

Next, I need to thank the people I have spent the past few years sitting next to everyday. Well, we didn't just sit there in our office. We also sat outside in Killian Court. We sat at Bonchon before they switched chicken distributors. We sat at each other's apartments, playing poker, celebrating birthdays. We sat in classrooms watching movies. Sometimes, we even got up and played some intramural soccer. To all the graduate students I have met in MSEAS (in alphabetical order), Abhinav Gupta, Aditya Ghodgaonkar, Aditya K. Saravanakumar, Alonso V. Rodriguez, Aman Jalan, Anantha Narayanan Suresh Babu, Chinmay Kulkarni, Clara Dahill, Corbin Foucart, Ellery Rajagopal, Jacob Heuss, Jing Lin, Kyprianos Gkirgkis, Manan Doshi, Manmeet Singh Bhabra, Marcos Tornielli, Michael Humara, Tony Ryu, Wael Hajj Ali, and Zach Duguid, thank you. Thank you for the countless research discussions and, far more importantly, for the friendship.

I'd like to single out those that have been here for my entire or a vast majority of my stay at MIT as I have gotten to know them the best. Wael took me under his wing and got me started in acoustics. Manny showed me the way to live with a light heart, always cheering me up with his infectious laughter. The three of us will share a bond forever as the acoustics boys. Manan may not be an acoustics boy, but he and I have a shared love for JAX and our newly acquired GPU. Though I can't say I share his love for open-source software, I certainly respect it. Abhinav was our fearless leader, telling Pierre we shouldn't work on weekends (ha!), and he planned the wildest MSEAS parties. Though Abhinav's absence could never be filled, Tony, also known as T-Bone, certainly tried as he took over as head of the party planning committee. Tony never let me down for the many favors I have asked; he may be too kind for his own good. Together with Corbin, we have been the traveling trio, attending conferences (vacationing) in San Diego and Berlin, both on the government's dime, and meeting up in Amsterdam too. Corbin has also served as my personal trainer/fitness instructor, motivating me to get moving physically, and reminding me to indulge in the joys of life outside of the lab. To all of you: I have truly cherished the bonds we made. I would not have made it through this program without each of you.

Finally, I owe an enormous debt of gratitude to my entire family. Through college and now grad school, I have been away from home for nearly eight years. Mom, Dad, Brian, Jason, Shaley, and Rachel, I miss you every day that I am away, and I hold dear the days we still get to enjoy together. Thank you for all of the texts, the phone calls, the FaceTimes, the Zooms, and the screen shares. Not every day is easy, but having you all as an unfaltering support system is the reason I can keep going and be my best.