High-order Discontinuous Galerkin Methods and Deep Reinforcement Learning with Application to Multiscale Ocean Modeling

by

Corbin Foucart

B.S. Stanford University (2015)
S.M., Massachusetts Institute of Technology (2019)

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

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2023

© Corbin Foucart 2023. All rights reserved.

The author hereby grants to MIT a nonexclusive, worldwide, irrevocable, royalty-free license to exercise any and all rights under copyright, including to reproduce, preserve, distribute and publicly display copies of the thesis, or release the thesis under an open-access license.

Author . . . .
Department of Mechanical Engineering
& Center for Computational Science and Engineering
August 2, 2023

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

Accepted by . . . .
Nicolas Hadjiconstantinou
Chairman, Department Committee on Graduate Theses

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

High-order Discontinuous Galerkin Methods and Deep Reinforcement Learning with Application to Multiscale Ocean Modeling

by

Corbin Foucart

Submitted to the Department of Mechanical Engineering
& Center for Computational Science and Engineering
on August 2, 2023, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

With the expanding availability of computational power, numerical modeling plays an increasingly pivotal role in the field of oceanography, enabling scientists to explore and understand ocean processes which are otherwise inaccessible or challenging to observe directly. It provides a crucial tool for investigating a range of phenomena from large-scale circulation patterns to small-scale turbulence, shaping our understanding of marine ecosystems, global climate, and weather patterns. However, this same wide range of spatiotemporal scales presents a distinct computational challenge in capturing physical interactions extending from the diffusive scale (millimeters, seconds) to planetary length scales spanning thousands of kilometers and time scales spanning millennia. Therefore, numerical and parameterization improvements have and will continue to define the state of the art in ocean modeling, in tandem with the integration of observational data and adaptive methods. As scientists strive to better understand multiscale ocean processes, the thirst for comprehensive simulations has proceeded apace with concomitant increases in computing power, and submesoscale resolutions where nonhydrostatic effects are important are progressively becoming approachable in ocean modeling. However, few realistic ocean circulation models presently have nonhydrostatic capability, and those that do overwhelmingly use low-order finite-difference and finite-volume methods, which are plagued by dispersive errors, and are arduous to utilize in general, especially on unstructured domains and in conjunction with adaptive numerical capabilities. High-order discontinuous Galerkin (DG) finite element methods (FEMs) allow for arbitrarily high-order solutions on unstructured meshes and often out-compete low-order models with respect to accuracy per computational cost, providing significant reduction of dispersion and dissipation errors over long-time integration horizons. These properties make DG-FEMs ideal for the next generation of ocean models, and, in this thesis, we develop a novel DG-FEM ocean model with the above longer-term vision and adaptive multiscale capabilities in mind.

Using a novel hybridizable discontinuous Galerkin (HDG) spatial discretization for both the hydrostatic and nonhydrostatic ocean equations with a free surface, we develop an accurate and efficient high-order finite element ocean model. We emphasize the stability and robustness properties of our schemes within a projection method discretization. We provide detailed benchmarking and performance comparisons for the parallelized implementation, tailored to the specifics of HDG finite element methods. We demonstrate that the model achieves optimal convergence, and is capable of accurately simulating nonhydrostatic behavior. We evaluate our simulations in diverse dynamical regimes including linear gravity waves, internal solitary waves, and the formation of Rayleigh-Taylor instabilities in the mixed layer. Motivated by investigating local nonhydrostatic submesoscale dynamics using realistic ocean simulation data, we develop schemes to initialize and nest the new DG-FEM model within a comprehensive hydrostatic ocean modeling system. Nested within such data-assimilative hydrostatic simulations in the Alboran Sea, we provide a demonstration of our new model's ability to capture both hydrostatic and nonhydrostatic dynamics that arise in the presence of wind-forced instabilities in the upper ocean layers. We show that such a model can both validate and work in tandem with larger hydrostatic modeling systems, enabling multi-dynamics simulations and enhancing the predictive fidelity of ocean forecasts.

Next, as DG-FEM methods are well-suited to adaptive refinement, we develop a method to learn new adaptive mesh refinement strategies directly from numerical simulation by formulating the adaptive mesh refinement (AMR) process as a reinforcement learning problem. Finite element discretizations of problems in computational physics can usefully rely on adaptive mesh refinement to preferentially resolve regions containing important features during simulation. However, most spatial refinement strategies are heuristic and rely on domain-specific knowledge or trial-and-error. We treat the process of adaptive mesh refinement as a local, sequential decision-making problem under incomplete information, formulating AMR as a partially observable Markov decision process. Using a deep reinforcement learning (DRL) approach, we train policy networks for AMR strategy directly from numerical simulation. The training process does not require an exact solution or a high-fidelity ground truth to the partial differential equation (PDE) at hand, nor does it require a pre-computed training dataset. The local nature of our deep reinforcement learning approach allows the policy network to be trained inexpensively on much smaller problems than those on which they are deployed, and the DRL-AMR learning process we devise is not specific to any particular PDE, problem dimension, or numerical discretization. The RL policy networks, trained on simple examples, can generalize to more complex problems and can flexibly incorporate diverse problem physics. To that end, we apply the method to a range of PDEs relevant to fluid and ocean processes, using a variety of high-order discontinuous Galerkin and hybridizable discontinuous Galerkin finite element discretizations. We show that the resultant learned policies are competitive with common AMR heuristics and strike a favorable balance between accuracy and cost such that they often lead to a higher accuracy per problem degree of freedom, and are effective across a wide class of PDEs and problems.

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

Acknowledgments

> The truth will set you free. But not until it is finished with you.
> — David Foster Wallace, *Infinite Jest*

There are a number of people who have been instrumental in my completion of this thesis, and to whom I am grateful.

I feel an enormous sense of gratitude to Pierre for his guidance throughout my time at MIT as well as his tireless work ethic. There's an apocryphal story I like about a wise and battle-hardened engineer called out of retirement as a consultant to fix a very expensive but malfunctioning machine. The engineer looks at the machine design and, after a minute or two of thought, draws with his pencil a small "x" on the part that needs to be replaced. After the replacement, the machine works perfectly. However, upon receiving a bill for half the original cost of the machine, the indignant vice president of the company demands an itemization. The engineer sends back a two-line invoice: fifty cents for the cost of the pencil and the rest for "knowing where to draw the x." Thank you, Pierre, for giving me the independence and space to work on my own ideas and mature as a researcher, but at the same time for always knowing exactly where the "x" should go.

I'd like to thank my committee for their comments and feedback. In particular, I'd like to thank Professor Peraire, Professor Patera, and Dr. Cuong Nguyen for meeting with me repeatedly over the course of my PhD to discuss research ideas and finite element methods more broadly. I'd like to thank Dr. Chris Mirabito for his near-infinite patience when I began my work on this project, not to mention his substantial contributions in proofreading the details of my work the past few years. Similarly, I owe a debt of gratitude to the legendary Dr. Pat Haley for more than living up to the title of "resident wise, old guy;" many of the ideas in these pages have directly benefited from your vast experience and bottomless well of debugging ideas. We are very grateful to the Office of Naval Research for support throughout this research and Ph.D. under grants N00014-15-1-2626 (DRI-FLEAT), N00014-18-1-2781 (DRI-CALYPSO), and N00014-20-1-2023 (MURI ML-SCOPE), the Defense Advanced Research Projects Agency (DARPA) for support under grant N66001-16-C-4003 (POSYDON), and the National Science Foundation for support under grants OCE-1061160 (ShelfIT) and EAR-1520825 (NSF-ALPHA), as well as to the Massachusetts Institute of Technology for making this research experience possible. Graduate school has taken me around the world and given me the opportunity to pursue anything and everything that I found worthwhile. I will always look back fondly on this period of my life.

Academically, I'm especially grateful to professors Boyd, Broderick, Marzouk, and Patera for their teaching—your ability to unpretentiously break down and distill complicated technical material to the point where it seems intuitive and easy has always impressed me. I'd also like to thank Dr. Alexander Linke of the Weierstrass Institute for his incredible mentorship during my time in Berlin—your zeal for applied math is infectious, and my choice to pursue research in computational physics was largely due to my time spent working with you. To Hui-Min, thank you for the years of advice, brutal honesty, friendship, and for pushing me, not only musically, but in nearly every aspect of life. You're the best mentor I've ever had and you've profoundly influenced me for the better.

To the friends I've made over the years at MSEAS, thank you for all the SGTs. Abhinav, thank you for being a continual grounding presence in my life and being the best friend I could possibly ask for. Aaron, I owe you a big "oh, danke" for the all the great vocalizations, animal noises, and memories traveling. To Arko, thank you for never failing to call me out when I believe something with too much certainty, for always knowing the shortest walking path between two points on campus, and for appreciating the nutritional importance of fried chicken. Manan, thank you for all the coding and statistical discussion over the years—on top of being an exceptional project partner, you may be the only other person I'll ever meet who appreciates Casella, Berger, and numerical software design as much as I do. To Manny, thank you for all the money you've given to me over the years in bets won; although you will always make me *so* mad, what you lack in wagering prowess you more than make up for in personality, music taste, and frisbee ability. Aditya, thank you for all your diligent work on the HDG project—it's rare to find someone who's as competent yet as easygoing as you are, and I will miss working with you. Tony, our spontaneous coffee trips have always been a welcome break from life's other obligations; may you continue to be loved by all people and feared by all geese. To Jing, your accipitrine attention to detail will always be an inspiration to me. Thank you for all the numerical intuition you gave me as a new PhD student and for the many enjoyable swim workouts in the "p-pool." And lastly, to JVo, I'll always look back fondly on our years working together; thanks for all the late nights spent debugging, drinking seltzer, and listening to Yuja Wang and Joshua Bell rip through the Kreutzer sonata, not to mention the great times in Miami, New England, Montreal, and not Albuquerque, where things are basically post-apocalyptic and where I would never visit you. Lisa, despite the fact that everyone loves you, you're still under-rated; thank you for all the admin work, laughs, and snacks. To my fellow night-owl John, thanks for keeping our building spotless and for the engaging 4AM conversations—I hope you enjoy a restful and well-deserved retirement! To the rest of you: thank you for the countless memories and for always making MSEAS feel like home.

I feel doubly fortunate to have made many friends during my time at MIT outside of academics. To Cindy, David, Emily, Florian, Zach, Haekyung, and Diana, thank you for the rewarding musical moments with Beethoven, Franck, Schubert, and Bach. To Tim, Misha, Peter and the other kick-boxers, thank you for pushing me physically, many times to the ground (ha), and perhaps for pushing me spiritually too. To Ben and Ved, I can't tell you how much fun it's been hitting the pavement with you two—thank you for nudging me out of my comfort zone. To Emma, Julian, Vivian, Nick, Harry, and Icey, thanks for the many coffee trips, regulatory compliance pain smiles, and laughs over the strange pastries and questionably-devised drinks of Kowloon Tong, TST, and LKF. I'd like to thank Emma in particular for being such a detail-oriented and patient language exchange partner; although I still don't think you'll ever look at the pictures you insisted on taking of every dog, flower, and piece of food we encountered, I'll always cherish our friendship. To Nayun, a profound thank you for helping me overcome my crippling fear of "being nice."

To Kevin, Evan, Michael William Locke, Nicky T., Richie, and Nate: words can't describe how great it's been growing up with all of you over the past 15 years. Through all of the many "camping" trips, poker and nighttime glow-in-the-dark ultimate frisbee games, themed parties, terrible movie nights, Blue Moon calzones, Castle Island runs, beach days, games of deception, weddings... although a lot has changed over the years, one thing stays the same—you guys are the best.

And, of course, I'd like to thank my amazing biological family, to whom I owe the gift of life itself. Thank you, Mom and Dad; everything I've accomplished has been in no small part due to the love, support, and guidance you've given me over the years. Thanks to my sister Abbey for keeping me centered and for all the excellent times at the Friendly Toast diner, and thanks to Esme for selflessly defending me from all the menacing air-conditioning units, autumn leaves, and stationary shadows of Boston, MA.