Workshop Report 23w5129
Scientific Machine Learning

Held June 18 to June 23 by the Banff International Research Station for Mathematical Innovation and Discovery (BIRS)

Brendan Keith[*] Thomas O'Leary-Roseberry[†] Lu Lu[‡]
Siddhartha Mishra[§] Zhiping Mao[¶]

November 14, 2023

Abstract

This report documents a 2023 Banff International Research Station for Mathematical Innovation and Discovery (BIRS) workshop that united international experts from academia, industry, and national laboratories in the field of Scientific Machine Learning (SciML). The workshop successfully fostered collaboration among researchers from diverse backgrounds, including applied mathematics, myriad engineering disciplines, scientific computing, optimization, and machine learning. The event endeavored to stimulate innovation in the field through scientific talks and open interactions aimed at various scientific and engineering applications.

1 Introduction

Scientific Machine Learning (SciML) is an evolving discipline that melds machine learning (ML) principles with scientific computing, paving the way for groundbreaking innovations in science and engineering. Recognizing the immense potential and rapid developments in this field, the 2023 BIRS Workshop on Scientific Machine Learning aimed to provide a fertile ground for collaboration and knowledge exchange.

2 Overview of the Field

SciML is a burgeoning interdisciplinary field that marries traditional scientific computing with advanced ML techniques. At its core, SciML seeks to harness the power of ML models, such as

[*]Division of Applied Mathematics, Brown University, Providence, RI 02912 USA (brendan_keith@brown.edu).
[†]Oden Institute for Computational Engineering & Sciences, The University of Texas at Austin (tom.olearyroseberry@utexas.edu).
[‡]Department of Statistics and Data Science, Yale University, New Haven, CT 06511 (lu.lu@yale.edu).
[§]Department of Mathematics, Rämistrasse 101, 8092 Zürich, Switzerland (siddhartha.mishra@sam.math.ethz.ch).
[¶]School of Mathematical Sciences, Xiamen University, Xiamen, 361005, China (zpmao@xmu.edu.cn)

deep neural networks, to accelerate and enhance the simulation, modeling, and analysis processes inherent in scientific disciplines. This fusion allows for more sophisticated handling of complex systems and phenomena, from predicting turbulent flows in fluid dynamics to optimizing molecular configurations in material science. Through leveraging data-driven ML methodologies alongside traditional physics-based models, SciML is redefining the boundaries of computational capabilities and the accuracy of predictions, making it possible to tackle previously intractable scientific problems.

However, the union of scientific computing and machine learning is not without its challenges. It demands a robust understanding of both domains, ensuring that the integrity of scientific principles is maintained while employing ML techniques. To this end, researchers in SciML are continuously working on innovative algorithms and frameworks that can seamlessly integrate the strengths of both worlds. The goal is to create models that are not only highly accurate but also interpretable, trustworthy, and aligned with the underlying scientific principles. As SciML continues to evolve, its applications are anticipated to proliferate across various scientific sectors, heralding a new era of research and discovery.

3 Recent Developments

SciML has witnessed many innovations in recent years, dramatically reshaping the landscape of scientific research. One of the most notable advancements is the integration of neural ordinary differential equations (neural ODEs). This approach allows researchers to leverage neural networks within ordinary differential equations, enabling the simultaneous learning of dynamic system behavior from data while maintaining a structured model format. This blend aids in improving computational efficiency, reducing the amount of data needed for accurate predictions, and providing a more interpretable model structure that aligns with scientific principles. Neural ODEs have found applications in various domains, including fluid dynamics, pharmacokinetics, and even in some areas of finance.

Another example of SciML are physics-informed neural networks (PINNs) [23]. These networks incorporate known physical laws (like conservation laws) through a loss function, effectively bridging the gap between data-driven approaches and theoretical knowledge. By constraining neural network training with these laws, PINNs ensure that the resultant models are not just data-adaptive but also physically consistent. This integration has proven particularly valuable in scenarios with sparse or noisy data, where traditional machine learning models might overfit or produce non-physical results.

A groundbreaking current research direction in SciML is *operator learning*. Operator learning broadly refers to various machine learning strategies for learning maps between (formally infinite-dimensional) function spaces, such as those containing functions that parametrize PDEs, as well as the functions who solve them. Accurate operator approximations can be substituted for expensive physics simulations in “many-query” settings which required repeated PDE solutions for differing parameters. Examples of many query problems include inverse problems (both deterministic and Bayesian), optimization problems (e.g., design and control) among other tasks. A major recent breakthrough in operator learning are the so-called “neural operators”, which use neural network approximations married with classical mathematical tools to exploit known structure of problems, while leveraging ML technologies to learn complex nonlinear representations from data and physics. For instance, Fourier Neural Operators [14] leverage the Fourier transform's efficiency to predict complex physical behaviors in various spatial dimensions. Another example is the DeepONet [15], which learns operators by treating them as mappings from functions to functions, utilizing two interacting neural networks — a branch network and a trunk network. The potential to represent and predict scientific phenomena with unprecedented accuracy is immensely expanded through these advanced neural architectures.

The 2023 BIRS Workshop on Scientific Machine Learning featured numerous talks on neural ODEs, PINNs, and neural operators, as well as numerous other emerging methodologies such as statistical inference methods and high performance computing considerations. The talks spanned algorithmic innovation, approximation theory and applications to complex physical systems.

4 Presentations

The following is a list of the presentations given at the workshop:

Monday, June 19, 2023

1. Panos Stinis, Pacific Northwest National Laboratory, Computational Mathematics Group: "Mutlifidelity Scientific Machine Learning". [10]

2. Romit Maulik, The Pennsylvania State University, Department of Information Science and Technology: "Multiscale Graph Neural Network Autoencoders for Interpretable Scientific Machine Learning" [2]

3. Scott Field, The University of Massachusetts, Dartmouth, Department of Mathematics: "Potential Applications of Scientific Machine Learning to the Binary Black Hole Problem" [12]

4. Bart van Bloemen Waanders, Sandia National Laboratories, Scientific Machine Learning group: "Learning control policies for high-fidelity models using hyper-differential sensitivities with respect to model discrepancy" [9]

5. Animashree Anandkumar, NVIDIA and California Institute of Technology, Computing and Mathematical Sciences: "Neural Operators for Accelerating Scientific Simulations " [14, 22]

6. Michael Brennan, Massachusetts Institute of Technology, Computational Science & Engineering: "Exploiting Low-Rank Conditional Structure to Solve Bayesian Inverse Problems" [4]

Tuesday, June 20, 2023

7. Yunan Yang, Cornell University, Department of Mathematics: "Neural Inverse Operators for Solving PDE Inverse Problems" [19]

8. Aras Bacho, Ludwig Maximilian University of Munich, Department of Mathematics: "PoissonNet: Resolution-Agnostic 3D Shape Reconstruction using Fourier Neural Operators" [1]

9. Eric Cyr, Sandia National Laboratories, Computational Mathematics Group: "Exploiting time-domain parallelism to accelerate neural network training" [20]

10. Paolo Zunino, MOX - Modelling and Scientific Computing - Politecnico di Milano: “A Deep Learning approach to Reduced Order Modeling of parameter dependent Partial Differential Equations” [6]

11. Robert Scheichl, Heidelberg University, Department of Mathematics: “Structure-preserving learning of embedded closure models for fluid flows”.

12. N. Sukumar, The University of California, Davis, Department of Civil and Environmental Engineering: “Exact Imposition of Boundary Conditions in PINNs to Solve PDEs”

13. Mihai Nica, the University of Guelph, Department of Mathematics and Statistics: “A derivative-free method for solving elliptic partial differential equations with deep neural networks” [8]

14. Deepanshu Verma, Emory University:, Department of Mathematics “Advances and challenges in solving high-dimensional Hamilton-Jacobi-Bellman equations”

Wednesday, June 21, 2023

15. Jakob Zech, Heidelberg University, Department of Mathematics: “Nonparametric Distribution Learning via Neural ODEs” [18]

16. Nicholas Nelsen, The California Insitute of Technology, Division of Engineering and Applied Science: “Convergence Theorem for Vector-Valued Random Features” [13]

17. Margaret Trautner, The California Insitute of Technology, Department of Computing + Mathematical Sciences: “Learning Homogenized Constitutive Laws” [3]

18. Guang Lin, Purdue University, Departments of Mathematics, Statistics & School of Mechanical Engineering: “ Energy Dissipative Evolutionary Deep Operator Networks” [24]

19. Thomas O’Leary-Roseberry, The University of Texas at Austin, Oden Institute for Computational Engineering & Sciences: “Derivative-Informed Neural Operators for High-Dimensional Outer-Loop Problems” [21]

20. Jinwoo Go, Georgia Institute of Technology, School of Computational Science and Engineering: “Accelerating A-Optimal/D-Optimal Design of Experiments Using Neural Networks”

21. Dingcheng Luo, The University of Texas at Austin, Oden Institute for Computational Engineering & Sciences: “Efficient PDE-constrained optimization with derivative-informed neural operators” [16]

22. Bruno Despres, Jacques-Louis Lions Laboratory Sorbonne University: “Generating functions for polynomials with ReLU: application to training”

23. Marta D’Elia, Pasteur Labs: “GNN-based physics solver for time-independent PDEs” [7]

24. Shunyuan Mao, University of Victoria, Department of Physics and Astronomy: “PPDONet: Deep Operator Networks for Fast Prediction of Steady-State Solutions in Disk-Planet Systems” [17]

Thursday, June 22, 2023

25. Petros Koumoutsakos, Harvard University, Institute for Applied Computational Science: “AI/Scientific Computing: Alloys for Flow modeling and Control” [11]

26. Peng Chen, Georgia Institute of Technology, School of Computational Science and Engineering: “ Projected variational inference for high-dimensional Bayesian inverse problems” [5]

27. Lu Lu, Yale University, Department of Statistics and Data Science: “Deep neural operators with reliable extrapolation for multiphysics, multiscale & multifidelity problems” [15, 25]

5 Scientific Progress Made

Participants of the 2023 BIRS Workshop on Scientific Machine Learning experienced a unique forum for knowledge exchange, collaborative problem solving, and exposure to cutting-edge theories and methods. The gathering served as a fertile ground for cross-pollinating ideas from various mathematical disciplines and applications while staying focused on developing new SciML methodologies. Senior researchers presented recent advancements, novel theories, and breakthrough methodologies, which helped to inspire younger researchers and offer fresh perspectives on both new and longstanding problems. Concurrently, the collaborative atmosphere of the workshop helped to catalyze the genesis of new research directions, as attendees often engaged in rigorous discussions that challenged existing paradigms and used the allotted time for unstructured interactions to brainstorm innovative solutions. Beyond the immediate dissemination of knowledge, we hope that the workshop has helped sow the seeds for future collaborations, publications, and even entirely new research domains, further solidifying BIRS’s role as an essential catalyst in the relentless journey of mathematical discovery.

6 Hybrid format

The hybrid format was beneficial because it allowed us to significantly expand the scope of our workshop. Due to the hybrid format, we were able to bring in a number of very high-impact remote presenters that would otherwise not have been able to speak. Additionally, we were able to engage a number of remote participants, including researchers who had difficulties obtaining Canadian visas in time and early career grad students who were not able to attend.

The major benefit of the workshop, however, was the in-person interactions and open discussions that happened in response to and in addition to the technical talks. Part of what really made the workshop work as well was the scale of it: since there were fewer participants, people got to know each other better than they otherwise would have, which led to deeper engagement with the technical material and enlightening discussions about the future of the field. For this reason, the hybrid format was non-essential, and it would make sense to make the event only in person.

7 Participant Testimonials

We received the following public testimonial from Robert Scheichl of the Institute for Mathematics & Interdisciplinary Center of Scientific Computing, Heidelberg University:

> "The workshop provided an excellent venue for interaction on this exciting and booming new research area. I would like to congratulate and thank the organisers for putting together a very well-balanced and high-profile line-up of talks. The interaction and engagement of participants at BIRS was great. It did showcase several new research avenues to me that will be highly influential in my future research; it did, however, also show the limitations of certain approaches. Participants were open for a non-competitive and unbiased discussion of competing potential approaches in the area of scientific machine learning. I particularly appreciated that a strong emphasis at this workshop was given to promising young and upcoming researchers, who brought a lot of enthusiasm, energy and willingness to engage and interact. The talks were well chosen and the things I took away will definitely influence my immediate hiring decisions. The concluding discussions on Friday on the wider discipline and on how this new field should be shaped was particularly useful and via an editorial piece we plan to write for a special journal issue should provide some lasting impact for the wider research field."

8 Outcomes of the Meeting

The workshop brought together an international and interdisciplinary group of researchers that would not otherwise have assembled to discuss the issues and challenges of the emerging field of scientific machine learning.

The countries represented were Canada, France, Germany, France, New Zealand, and the USA. The disciplines represented include computational and applied mathematics, information science, scientific machine learning, computational science & engineering, aerospace engineering, chemical engineering, civil engineering, mechanical engineering, physics, and astrophysics, spanning academia, labs, and industry.

In the workshop, key emerging areas of promise within SciML were identified, while honest critical debate led to a better understanding of the limitations of the field. This gave clarity to future directions of interest in the field. We hope that numerous collaborations and international research correspondences will arise due to this workshop. The organizers can attest personally that this is already a takeaway from this event.

The workshop will culminate in a special issue (SI) in the American Institute for Mathematical Sciences Foundations of Data Science (FoDS) journal. The SI will be guest-edited by the BIRS organizers. Additionally, a position paper summarizing the findings of some discussions at the workshop will accompany this issue.

References

[1] Hector Andrade-Loarca, Julius Hege, Aras Bacho, and Gitta Kutyniok. Poissonnet: Resolution-agnostic 3d shape reconstruction using fourier neural operators. *arXiv preprint arXiv:2308.01766*, 2023.

[2] Shivam Barwey, Varun Shankar, and Romit Maulik. Multiscale graph neural network autoencoders for interpretable scientific machine learning. *arXiv preprint arXiv:2302.06186*, 2023.

[3] Kaushik Bhattacharya, Nikola Kovachki, Aakila Rajan, Andrew M Stuart, and Margaret Trautner. Learning homogenization for elliptic operators. *arXiv preprint arXiv:2306.12006*, 2023.

[4] Michael Brennan, Daniele Bigoni, Olivier Zahm, Alessio Spantini, and Youssef Marzouk. Greedy inference with structure-exploiting lazy maps. *Advances in Neural Information Processing Systems*, 33:8330–8342, 2020.

[5] Peng Chen, Keyi Wu, Joshua Chen, Thomas O'Leary-Roseberry, and Omar Ghattas. Projected Stein variational Newton: A fast and scalable Bayesian inference method in high dimensions. *Advances in Neural Information Processing Systems*, 32, 2019.

[6] Nicola Franco, Andrea Manzoni, and Paolo Zunino. A deep learning approach to reduced order modelling of parameter dependent partial differential equations. *Mathematics of Computation*, 92(340):483–524, 2023.

[7] Rini Jasmine Gladstone, Helia Rahmani, Vishvas Suryakumar, Hadi Meidani, Marta D'Elia, and Ahmad Zareei. Gnn-based physics solver for time-independent pdes. *arXiv preprint arXiv:2303.15681*, 2023.

[8] Jihun Han, Mihai Nica, and Adam R Stinchcombe. A derivative-free method for solving elliptic partial differential equations with deep neural networks. *Journal of Computational Physics*, 419:109672, 2020.

[9] Joseph Hart, Bart van Bloemen Waanders, Lisa Hood, and Julie Parish. Sensitivity-driven experimental design to facilitate control of dynamical systems. *Journal of Optimization Theory and Applications*, 196(3):855–881, 2023.

[10] Amanda A Howard, Mauro Perego, George E Karniadakis, and Panos Stinis. Multifidelity deep operator networks. *arXiv preprint arXiv:2204.09157*, 2022.

[11] Petr Karnakov, Sergey Litvinov, and Petros Koumoutsakos. Optimizing a discrete loss (odil) to solve forward and inverse problems for partial differential equations using machine learning tools. *arXiv preprint arXiv:2205.04611*, 2022.

[12] Brendan Keith, Akshay Khadse, and Scott E Field. Learning orbital dynamics of binary black hole systems from gravitational wave measurements. *Physical Review Research*, 3(4):043101, 2021.

[13] Samuel Lanthaler and Nicholas H Nelsen. Error bounds for learning with vector-valued random features. *arXiv preprint arXiv:2305.17170*, 2023.

[14] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *International Conference on Learning Representations*, 2021.

[15] Lu Lu, Pengzhan Jin, Guofei Pang, and George Em Karniadakis. DeepONet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 2021.

[16] Dingcheng Luo, Thomas O'Leary-Roseberry, Peng Chen, and Omar Ghattas. Efficient pde-constrained optimization under high-dimensional uncertainty using derivative-informed neural operators. *arXiv preprint arXiv:2305.20053*, 2023.

[17] Shunyuan Mao, Ruobing Dong, Lu Lu, Kwang Moo Yi, Sifan Wang, and Paris Perdikaris. Ppdonet: Deep operator networks for fast prediction of steady-state solutions in disk–planet systems. *The Astrophysical Journal Letters*, 950(2):L12, 2023.

[18] Youssef Marzouk, Zhi Ren, Sven Wang, and Jakob Zech. Distribution learning via neural differential equations: a nonparametric statistical perspective. *arXiv preprint arXiv:2309.01043*, 2023.

[19] Roberto Molinaro, Yunan Yang, Björn Engquist, and Siddhartha Mishra. Neural inverse operators for solving pde inverse problems. *arXiv preprint arXiv:2301.11167*, 2023.

[20] Gordon Euhyun Moon and Eric C Cyr. Parallel training of gru networks with a multi-grid solver for long sequences. *arXiv preprint arXiv:2203.04738*, 2022.

[21] Thomas O'Leary-Roseberry, Peng Chen, Umberto Villa, and Omar Ghattas. Derivative-Informed Neural Operator: An efficient framework for high-dimensional parametric derivative learning. *arXiv preprint arXiv:2206.10745*, 2022.

[22] Jaideep Pathak, Shashank Subramanian, Peter Harrington, Sanjeev Raja, Ashesh Chattopadhyay, Morteza Mardani, Thorsten Kurth, David Hall, Zongyi Li, Kamyar Azizzadenesheli, et al. Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. *arXiv preprint arXiv:2202.11214*, 2022.

[23] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.

[24] Jiahao Zhang, Shiheng Zhang, Jie Shen, and Guang Lin. Energy-dissipative evolutionary deep operator neural networks. *arXiv preprint arXiv:2306.06281*, 2023.

[25] Min Zhu, Handi Zhang, Anran Jiao, George Em Karniadakis, and Lu Lu. Reliable extrapolation of deep neural operators informed by physics or sparse observations. *Computer Methods in Applied Mechanics and Engineering*, 412:116064, 2023.