Structured Machine Learning and TimeStepping for Dynamical Systems (24w5301)

Jason Bramburger (Concordia), Chris Budd (UBath),
Jingwei Hu (UWashington), Andy Wan (UC Merced)

February 19-23, 2024

1 Overview of the Field

Machine learning, in particular Deep Learning is one of the fastest growing areas of modern science and technology, which has potentially enormous and transformative impacts on all areas of our life. However, for many physical problems where Deep Learning is currently being applied, one distinguishing feature is the plethora of domain knowledge and historical data accumulated over decades, yet with scarcity in labelled data. This currently does not fit well within the supervised machine learning paradigm, which is mostly agnostic to physical models and requires a large amount of training data in order to make accurate and generalisable predictions. In particular, while the supervised machine learning approach possesses superior ability for interpolation, it has significant shortcomings when used for extrapolation, in contrast to models based on physics or possessing inherent structures. Thus, there is a great interest and emergent need to develop the next generation of machine learning algorithms guided by the underlying physics or structures that is capable of combining established domain knowledge and powerful machine learning techniques to make accurate and reliable predictions with limited data. This new paradigm of *structured machine learning* has the potential to significantly improve the efficiency of conventional algorithms and to greatly accelerate scientific investigations, with much better interpretability than purely data-driven generative machine learning algorithms.

To improve the accuracy of predictions for dynamical systems, the idea of incorporating inherent structures into numerical algorithms has been investigated thoroughly in recent decades [14]. Specifically, structure-preserving time-stepping algorithms such as symplectic integrators, variational integrators, conservative integrators, Lie group integrators, splitting methods, IMEX methods, strong-stability-preserving methods, asymptotic-preserving methods have all been applied successfully to find accurate approximations for a wide variety of dynamical systems or differential equations arising from the mathematical sciences and engineering. Such methods aim to preserve the inherent structures of differential equations, such as conservation laws, dissipative properties, symmetries, positivity, asymptotic limits, and more. For instance, for Hamiltonian systems, it has been well understood that symplectic integrators can provide more stable and accurate numerical

results in long term simulations over traditional integrators [14, 23]. For dynamical systems with conserved quantities, it has also been shown that exactly conservative integrators [24] can provide more stable numerical solutions for an arbitrarily long time [25] over traditional and symplectic integrators. For hyperbolic conservation laws satisfying a maximum principle, it is known that strong-stability-preserving methods [13] are the ideal framework to preserve any convex functional bound. Moreover, for dissipative systems, such as kinetic equations or gradient flows endowed with an entropy or a free energy functional, special time-stepping methods can be designed to preserve the entropy or energy decay structure; see for example [15, 1, 12]. These successes from numerical analysis naturally motivate the central question of this meeting: *can similar ideas can be incorporated into machine learning algorithms and architecture design of deep learning models*? If so, then the potential benefits are enormous in the design, analysis, and application of such methods which can enable accurate and reproducible predictions for dynamical systems.

Simultaneously with the above, advances in machine learning and deep neural networks over the past decade have enabled computational scientists to solve ever-increasingly complex problems in dynamical systems. This includes problems that have gone back decades, such as identifying nonlinear projections of spatio-temporal chaos onto low-dimensional inertial manifolds [16], identifying global coordinate transformations that linearize complex systems [19], and learning nonlinear reduced-order models [4, 9]. More recently, various connections between deep neural network architectures and time-stepping algorithms have also been identified, such as in the design of stable neural network architectures [8], and accelerated gradient descent and sampling algorithms [2] inspired by structure-preserving time-stepping methods. Currently, machine learning techniques used for predicting the evolution of a dynamical system typically employ generative algorithms trained on observed, or calculated data, such as in the use of parameterizations in weather forecasting. Such an approach can be characterized by learning a nonlinear evolution operator used to simulate the underlying dynamics. However, there is no theoretical guarantee that the resulting operator can well approximate the underlying evolution operator of the dynamical system, especially when the learned model is used for prediction at regions of the phase space distant from the training data. This leads to a fundamental question of *how well the learned operator corresponds to the underlying evolution operator of the dynamical system*. This question has already been posed in the context of numerical analysis [23] and it is well-known that structure-preserving time-stepping methods, which preserve the structure of conservation laws, symmetries, or other geometric features of the dynamical system, can provide more accurate and stable predictions over long term simulation than traditional numerical methods based on discretisation error alone.

In addition to structure-preserving time-stepping algorithms, incorporating structures directly into machine learning algorithms also possesses great potentials at improving accuracy and efficiency in approximating the evolution of dynamical systems. Specifically for evolutionary partial differential equations, fine discretization is often required to accurately resolve all relevant spatial and temporal scales in order to produce a high-fidelity representation of the dynamics, which unfortunately is responsible for the prohibitively expensive computations in traditional time-stepping algorithms. The need to resolve the underlying physics on both fast time scales and fine spatial scales places severe practical limitations on simulation times and the ability to model slow timescale and micro-scale physical processes. Traditional time-stepping schemes are typically designed using Taylor expansions, which are local in time and space with their accuracy determined by their step sizes. In contrast, there is a growing effort in the developing deep neural networks which can learn time-stepping and spatial-resolving schemes that are unconstrained by local Taylor expansions [20, 22, 21]. In this direction of machine learning based time-stepping schemes, recent efforts have focused on the viewpoint of flow maps in dynamical systems, with the objective of learning such maps that are unrestricted by the locality of Taylor expansions [6]. These flow maps can also

encode hierarchical time-stepping schemes to better capture the structure of the underlying flow maps over a disparate range of time-scales, with recent results suggesting their capability of overcoming traditional numerical stiffness [17]. However, almost no rigorous analysis has been shown to guarantee the performance of such algorithms. On the other hand, nearly every sub-field of numerical integration can be revisited from this perspective, with potentially even more efficient and accurate time-stepping schemes if adaptive step sizes are incorporated into this framework. Moreover, recent advances in capturing multiscale spatio-temporal resolutions have been made by deep convolutional autoencoders which implements numerical integration using ideas from multigrid methods and transfer learning [18]. Despite their demonstrated successes, there is currently a lack of theoretical understanding and rigorous error bounds with these approaches. Consequently, this recent line of research naturally motivates a key question on *how to leverage the wealth of numerical analysis knowledge to provide rigorous estimates on stability and accuracy for such structured deep learning based approaches*.

Thus, with the emergent interests at utilizing machine learning techniques for predicting complex dynamics, along with the relatively mature field of structure-preserving time-stepping algorithms, and the recent advances in deep learning approaches based on flow maps and multiscale resolution, we are now at a critical junction to bring together researchers from these distinct fields, whom would otherwise rarely interact with each other. It is therefore a central goal of this workshop to combine experts from these areas to devise a new paradigm of structured machine learning techniques for dynamical systems, inspired by structure-preserving time-stepping algorithms and the structures of flow maps and multiscale resolution.

This workshop aimed to bring together experts in both machine learning and numerical analysis in an effort to combine the flexibility and expressive power of neural networks, with the benefits of structure-preserving numerical methods and their rigorous analysis. The fusion of these ideas was exciting and mutually beneficial, as we spell out in the remaining sections of this report. We believe that results stemming from this workshop will have profound impacts at devising machine learning algorithms which are reliable, trustworthy, and flexible in major applications throughout the physical sciences.

2 Recent Developments and Open Problems

The integration of numerical analysis, dynamical systems and machine learning is a relatively new development in mathematics, with many contributions coming within the past decade. Many of the foundational contributions have been experimental, aiming to see what can be achieved and how far it can be pushed. While we now have a wealth of experimental evidence to draw upon, what is now required is proper mathematical analysis that confirms these experimental findings with explicit error bounds and performance guarantees.

In what follows we will highlight the major recent developments in the field, as well as outline some open problems related to them. Many of the related questions are the result of the open problem session of the workshop, held on February 22, 2024. An overarching theme of the discussion by the participants was to achieve a greater reliability of machine learning methods for time-stepping dynamical systems. In particular, we wish to determine going forward whether machine learning can beat (or is even competitive with) traditional time-stepping methods that have been developed by numerical analysts over the past half-century or more.

Neural Differential Equations. Lying exactly at the intersection of time-stepping dynamical systems and machine learning are neural differential equations [10] of the form $dx/dt = f(x, t)$ where

$f(x,t)$ is a the output of a neural net. These network architectures take advantage of the expressive capability of neural networks to identify the right-hand-side of a dynamical system. Unlike earlier methods of model identification, neural differential equations have the advantage that they do not rely on basis/library expansions and do not require estimations of the derivative from continuous-time processes. Nor do they need equally spaced data points for their training. Precisely, the network attempts to fit a finite collection of datapoints to a differential equation by attempting to identify the *exact* right-hand-side by minimizing a loss function that time-steps the differential equation using a traditional method from numerical analysis. Neural differential equations and their various offshoots were the topic of the first day of our workshop.

Being a recent innovation, there are many important questions that remain outstanding for neural differential equations. Some are related to the specific network architectures, such as which activation functions are best-suited to certain tasks/data (observe that if a ReLU activation is used then a Neural ODE has a non-smooth right hand side), or how sensitive the models are to hyperparameters and the data used to seed the method. For example, can performance and error bounds be determined in terms of the sampling rate of the training data? Furthermore, what is the effect of noisy measurements gathered from real-world as opposed to synthetic data and how can we build models that are robust to such corruption? At the mathematical level, we can ask how traditional time-stepping methods interface with neural differential equations in the loss function to determine how they can lead to a better understanding of error control. Finally, a theme of many presentations of this workshop focused on incorporating structure such as conserved quantities and symmetries into machine learning methods such as neural differential equations. It remains to determine if exploiting structure produces better models overall and whether it makes the neural network harder to train.

Physics-Informed Neural Networks. Innovations in machine learning such as back propagation have led researchers to use neural networks to approximate solutions to partial differential equations (PDEs). The idea is to take advantage of the same automatic differentiation used to train the neural network to differentiate its output with respect to its inputs, which are the (continuous) space and time independent variables of a PDE. Physics-informed neural networks (PINNs) get their name from the loss function representing the 'physics' of the model, here meaning it trains the output of the network to fit the PDE using the space and time differentials. PINNS come in two forms. The most commonly used methods are those which which evaluate a PDE residual based loss function at a set of collocation points (and hence are a form of collocation methods). Alternatively PINNS (such as the Deep Ritz method) which are based on the calculus of variations, in which casss the loss function is related to the variational functional. Partly due to their ease of use, and ability to work in very high dimensions, PINNs have been used prolifically, spawning at least three review papers in the six years since their inception [3, 7, 11].

Despite the explosion of interest in PINNs, it is still not clear whether they can outperform traditional methods of numerical analysis. This primarily comes from the lack of rigorous convergence theory for PINNs, meaning at present we only have experimental heuristic evidence for their utility. Moving forward it is necessary to establish whether PINNs can be competitive with traditional time-stepping methods of PDE simulation, including establishing their numerical stability. It would further be of benefit to the community to establish a database of benchmark problems that can be used to compare with more established methods. This would also help researchers to identify which models PINNs fail on, with the goal of finding out why, and also what sort of problems are well solved by a PINN which cannot be solved by a more traditional numerical method.

The Koopman Operator Framework. Despite being introduced in the 1930s, B.O. Koopman's

dynamical systems framework through the composition operator that now bears his name is drawing significant mathematical interest. The Koopman operator framework presents an equivalent linear formulation of nonlinear systems by lifting the dynamics to an infinite-dimensional Banach space of scalar functions called *observables*. The advantage of the linear Koopman description of the dynamics is that its temporal evolution can be completely understood in terms of its spectrum. Importantly, Koopman eigenfunction expansions constitute a space-time separation of variables that can be used to identify attractors, coherent sets, conserved quantities, and construct Lyapunov functions. The emergence of the Koopman operator again in the 21st century now comes from our ability to approximate it directly from data, providing a linearized description of genuinely nonlinear dynamics [19].

Much like PINNs, the Koopman operator framework has exploded in popularity in recent years. Unlike PINNs, there is significant rigorous theory that can justify its performance, with researchers now seeking to identify just how far the framework can be pushed to tackle increasingly complex problems. For example, can the Koopman framework be used to produce reduced-order models on low-dimensional manifolds for high-dimensional data? How can we incorporate structure to better inform learned Koopman models? Can we use the learned approximation of the Koopman operator extrapolate beyond the training set. Indeed, can the Koopman operator be applied to describe qualitatively different types of dynamics arising from the same (nonlinear) dynamical system (a good example is the motion of the simple pendulum). For this question there is some recent progress indicating that this is the case [5], but much work remains to be done both theoretically and experimentally. Can we reliably approximate chaotic dynamics within the Koopman framework? Finally, is there a meaningful connection between the Koopman operator framework and topological data analysis, and if so, how can this be used to further our understanding of nonlinear systems?

While the previous three points were targeted at dynamical systems, the following two points are aimed at the field of machine learning in general. The following open problems apply broadly to computational methods such as optimization, sampling, and neural networks. While these topics loom large over the field of machine learning, here we only target the discussion towards the theme of this workshop.

Optimization and Sampling. Underlying all machine learning algorithms is optimization and data sampling. Data-driven methods rely on minimizing loss functions, which are primarily constructed in terms of the data used to seed the method. Thus, everything is highly dependent on how the loss function is minimized, which loss function is used, and what data is provided to train the method. This therefore results in problems that underscore all machine learning methods, not just those related to time-stepping dynamical systems. Precisely, what are the best loss functions and regularizers to use in training neural network models and what is the optimal way to sample data for learning dynamical features of a system?

Neural Network Expressivity. Neural networks as a computational tool rely on theory that goes back decades into dense subsets of function spaces. These results are known as *universal approximation theorems* and have formed the backbone that modern machine learning now relies upon. A question that has arisen at the workshop is whether we can develop these theorems independently of traditional approximation theory. If not, this means that our theory can only be as good as so-called classical methods, such as those that underscore much of modern numerical analysis. On the computational side, while much of the theory has now been laid down, what remains are practical implementations of it with relatively low computational cost. For example, how can we construct neural networks that are effectively trainable to achieve the expressivity required for a given task?

3 Presentation Highlights

There were a total of 27 talks in our workshop, with 25 in-person talks and 2 virtual talks. Each workshop day was scheduled with topics related to our workshop themes, as we summarize next.
Day 1: Monday February 19. *Neural ODEs, generative diffusion model, learning conserved quantities and Lagrangian, stability and practical training of neural networks*

- *Subgrid-Scale Operators with neural ODEs*
  Emil Constantinescu (Argonne National Laboratory)
- *Numerical integrators for learning neural ordinary differential equation models*
  Takaharu Yaguchi (Kobe University)
- *Dynamical systems in deep generative modelling*
  Lisa Kreusser (University of Bath)
- *Machine learning of conservation laws for dynamical systems*
  HongKun Zhang (University of Massachusetts Amherst)
- *Learning Lagrangian dynamics from data with UQ*
  Christian Offen (Paderborn University)
- *Stability of numerical methods on Euclidean spaces and manifolds with applications to NNs*
  Brynjulf Owren (Norwegian University of Science and Technology)
- *Practical existence theorems for deep learning approximation in high dimensions*
  Simone Brugiapaglia (Concordia University)

There were 7 in-person talks on the first day. First, Emil Constantinescu highlighted how low fidelity source terms can be combined with a high fidelity source terms up-scaled via a neural ODE approach. Takaharu Yaguchi introduced stability concepts for Neural ODEs and showed theoretical limitations of learning vector field with Runge-Kutta integrators. Lisa Kreusser reviewed the score based diffusion model and showed that the error measured in Wasserstein 2-distance between the ODE and SDE induced distributions can be bounded above by a Fokker-Planck residual. HongKun Zhang discussed the Neural Deflation approach to learn additional functionally independent conserved quantities for integrable Hamiltonian systems such as Calogero-Moser system. Christian Offen introduced a Gaussian Process framework to learn a unique Lagrangian for variational systems from data by introducing gauge constraints and nondegeneracy condition. Brynjulf Orwen gave a new notion of B-stability on Riemannian Manifolds and as a result showed surprising results on non-uniqueness of geodesic implicit integrators on positively curved spaces. Finally, Simone Brugiapaglia presented practical existence results to bridge gaps between theory of universal approximation and practical training strategy by sparse high-dimensional polynomial approximation.

Day 2: Tuesday February 20. *Graph Neural Network, Geometric Machine Learning, Sampling, Optimization on non-Euclidean spaces, and connections with information geometry*

- *Improving the robustness of Graph Neural Networks with coupled dynamical systems*
  Davide Murari (Norwegian University of Science and Technology)
- *Time dependent graph neural networks*
  Eldad Haber (University of British Columbia)
- *Representation Trade-Offs in Geometric Machine Learning*
  Melanie Weber (Harvard University)

- *Conservative Hamiltonian Monte Carlo*
  Geoffrey McGregor (University of Toronto)
- *(Lie-group) Structured Inverse-free Second-order Optimization for Large Neural Nets*
  Wu Lin (Vector Institute)
- *Optimization and Sampling in Non-Euclidean Spaces*
  Molei Tao (Georgia Institute of Technology)
- *The Connections Between Discrete Geometric Mechanics, Information Geometry, Accelerated Optimization and Machine Learning*
  Melvin Leok (University of California San Diego)

The second day consisted of 6 in-person talks and 1 virtual talk which began with Davide Murari discussing the use of contractive dynamical systems to improve robustness against adversarial perturbations on Graphical Neural Network. This was followed up with Eldad Haber who introduced Differential Equation-Inspired Graph Neural Networks and highlighted the importance of using neural network architectures that mimic the underlying PDE properties. Delivered virtually, Melanie Weber reviewed the correlational statistical query model to illustrate the hardness of learning neural network and how incorporating data geometry can alleviate some but not all of the exponential complexity. Geoffrey McGregor introduced an extension of Hamiltonian Monte Carlo using conservative integrators and showed improvements in convergence and robustness for some high dimensional sampling problems. Wu Lin introduced a structured inverse-free second-order optimization called natural-gradient descent method on Gaussian manifold which is computationally and memory efficient. Motivated by optimization on Stiefel manifold, Molei Tao discussed variational Riemannian optimization and improving sampling from distribution on Lie groups via momentum sampler. Finally, Melvin Leok discussed the close connections between discrete geometric mechanics, information geometry and accelerated optimization.

Day 3: Wednesday February 21. *Shape optimization, control and uncertainty quantification for plasma physics, Learning Koopman operators for control.*

- *Deep neural networks on diffeomorphism groups for optimal shape reparameterization*
  Elena Celledoni (Norwegian University of Science and Technology)
- *Control and neural network uncertainty quantification for plasma simulation*
  Giacomo Dimarco (University of Ferrara)
- *Computationally Efficient Data-Driven Discovery and Linear Representation of Nonlinear Systems For Control*
  Bethany Lusch (Argonne National Laboratory)

The third day was a short day consisting with 3 in-person talks, starting with Elena Celledoni discussing the use of shape analysis to improve classification of curves, surfaces and manifolds with applications to computer vision. Then, Giacomo Dimarco introduced the control and uncertainty quantification problem for plasma simulations, as well as the use of Physics-Informed Neural Network as low fidelity models to improve sampling errors. Lastly, Bethany Lusch reviewed Koopman mode decomposition and how to represent nonlinear control dynamics via variational autoencoders and linear Koopman operators in latent spaces.

Day 4: Thursday February 22. *Expressiveness of ReLU networks, efficient learning with multiscale data, learning low-dimensional representation of chaotic systems, and connections with traditional numerical methods.*

- *Adaptivity and expressivity in neural network approximations*
  Chris Budd (University of Bath)
- *Efficient gradient descent algorithms for learning from multiscale data*
  Yen-Hsi Tsai (University of Texas Austin)
- *Data-driven modeling of complex chaotic dynamics on invariant manifolds*
  Michael Graham (University of Wisconsin Madison)
- *A spatiotemporal discretization for diffeomorphism approximation*
  Seth Taylor (McGill University)
- *Limited area weather modelling: interpolating with neural networks*
  James Jackaman (Norwegian University of Science and Technology)
- *A particle method based on Voronoi decomposition for the CahnHilliard equation*
  Daisuke Furihata (Osaka University)
- *Explicit time discretizations that preserve dissipative or conservative energy dynamics*
  David Ketchensen (King Abdullah University of Science and Technology)

The fourth day featured 7 in-person talks, starting with Chris Budd connecting Free Knot Splines approximation with the expressivity of shallow RELU neural networks, which provided insights on improving training of such networks. Yen-Hsi Tsai utilized multiscale expansion of loss functions and showcased examples on how to improve training via a multiscale variant of stochastic gradient descent. Michael Graham employed Data-driven Manifold Dynamics to learn invariant stable manifolds which can dramatically reduce degrees of freedom needed on examples such as KuramotoSivashinsky equation and Couette Flow. Seth Taylor reviewed the Characteristic Mapping Method by leveraging compositional nature of flow maps, leading to high fidelity solutions which mimics energy cascade for barotropic flows on the sphere. James Jackaman discussed solving limited area weather modelling by using graphical neural network to upscale low fidelity solutions to finer grid. Daisuke Furihara reviewed the Discrete Variational Derivative Method for the Cahn-Hilliard equation using a particle method and a dual grid based on Voronoi decomposition. Finally, David Ketchensen reviewed the relaxation Runge-Kutta method and showcased its effectiveness for various conservative and dissipative dynamics.

Day 5: Friday February 23. *Manifold learning, data driven discovery of invariant PDEs, and learning hyperbolic moment closure for kinetic equations.*

- *Manifold Flows and Generative Models*
  Kyriakos Flouris (ETH Zrich)
- *Learning PDEs from image data using invariant features*
  Yolanne Lee (University College London)
- *Hyperbolic machine learning moment closures for kinetic equations*
  Juntao Huang (Texas Tech University)

The last day was also a short day consisting of 2 in-person talks and 1 virtual talk. Kyriakos Flouris reviewed Normalizing flows and Manifold Learning and introduced Canonical Manifold Learning flows to improve performance with intrinsic dimension estimation and cosine similarity metrics. Yolanne Lee introduced learning PDEs with invariant polynomial differential operators to improve robustness against data transformation and reduce model complexity. Finally, delivered virtually, Juntao Huang reviewed the hyperbolic moment closure problem and introduced a gradient-based moment closure method to enforce hyperbolicity of neural networks approximating the higher order moment closure.

4 Scientific Progress Made

This meeting brought together experts in numerical analysis, machine learning, dynamical systems, and application areas. The result was a melding of communities with the intention, and action, of developing theories and computational techniques that can integrate machine learning with dynamical systems theory and numerical analysis. In line with the stated purpose of the meeting, this is a two way interaction. It is now very clear machine learning can be used to identify and simulate dynamics, and research needs to be done on identifying the strengths and limitations of this approach. Similarly, it became much clearer how ideas from dynamics and numerical analysis (such as optimisation, Neural ODEs, Neural Operators, DMD, etc.) can be used to inform, develop, explain (for example error analysis), and improve, machine learning algorithms. The relevance of these ideas to such diverse application areas as plasma physics, environmental sciences, and physiology.

Through their participation in this workshop, the invitees helped to solidify many of the open problems that are necessary to developing fast and reliable machine learning methods to better understand time-dependent systems. The alignment of the community came first from members highlighting their state-of-the-art research into both computational and analytical methods for interpreting, forecasting, and extracting information from temporally evolving equations and data. Towards the end of the meeting we held an open problems session for the participants to unify approaches moving forward. Many of these open problems are summarized in Section 2 above.

To maintain unified development, and accelerate progress, in the field of machine learning for time-stepping dynamical systems, and motivated by the panel discussion at the workshop, the participants and organizers have decided on the following measures moving forward:

1. Construction of a database of benchmark challenge problems (both theoretical and drawn from applications). These problems will be used to gauge performance of new developments in the field and compare them to traditional methods of numerical analysis. The database will be shared widely with other communities engaged in all aspects of the development of scientific machine learning.

2. The development of an online portal where members can find interested collaborators from the meeting. We envision this continue the unification of expertise brought together by this meeting and guarantee that participants can easily be connected with those who have the necessary skill-set that can further their research program.

3. The organizers will continue to organize annual or bi-annual check-ins with the participants through conferences and workshops. This will provide a mechanism for the quick dissemination of recent advances in the field, while also providing the opportunity for collaboration and further identification of open problems. In doing so they will collaborate with other initiatives, such as REMODEL (which involves many of the participants of this workshop) and the Common Task Framework, organised by Prof Nathan Kutz, also engaged with scientific machine learning lined to dynamical systems and numerical analysis.

To summarize, this workshop provided the necessary initial steps to unify and advance the research programs of mathematicians in the rapidly growing area of scientific machine learning, linked to dynamical systems theory and numerical analysis. We have identified many new problems to work on going forward, important theoretical and applicationss based challenges, and who is best suited to work on each one. It will provide for the opportunity to disseminate new results in the near future.

5 Outcome of the Meeting

The workshop drew speakers from various career stages, ranging from full professors to associate and assistant professors, postdoctoral researchers, and senior graduate students. Moreover, it also attracted speakers from national laboratories and companies, creating a diverse and dynamic lineup of experts. This blend of academic, industry, and research professionals ensures that attendees received insights from a wide range of perspectives and backgrounds, enriching the overall experience of the workshop. Additionally, the workshop prioritized gender balance, ensuring representation from both male and female researchers across various career stages, further enhancing the diversity and inclusivity of the event.

While most of the talks were held in person, each day, we had a designated chairperson to monitor the Zoom session, ensuring seamless communication between the in-person audience and the virtual audience. This setup allowed for full participation and engagement from both physical and remote attendees, creating an inclusive and interactive experience for all participants.

The highlight of the workshop was the open problem session held on Thursday afternoon. This interactive session provided a platform for attendees to engage in lively discussions, brainstorming solutions, and sharing innovative ideas. Some of the discussions were summarized in Section 2, and conclusions/next steps from them in Section 4. Following the workshop, a Google Sheet containing the contact information of all in-person participants, as well as notes from the open problems/discussion session, was shared with all participants. Many people have signed up for the topics they are interested in, and potential collaborations are likely to continue, fostering further exchange and advancement in the field.

References

[1] Rafael Bailo, José A Carrillo, and Jingwei Hu. Fully discrete positivity-preserving and energy-dissipating schemes for aggregation-diffusion equations with a gradient-flow structure. *Communications in Mathematical Sciences*, 18(5):1259–1303, 2020.

[2] Alessandro Barp, Lancelot Da Costa, Guilherme França, Karl Friston, Mark Girolami, Michael I Jordan, and Grigorios A Pavliotis. Geometric methods for sampling, optimization, inference, and adaptive agents. In *Handbook of Statistics*, volume 46, pages 21–78. Elsevier, 2022.

[3] Jan Blechschmidt and Oliver G Ernst. Three ways to solve partial differential equations with neural networks – a review. *GAMM-Mitteilungen*, 44(2):e202100006, 2021.

[4] Jason J Bramburger, Steven L Brunton, and J Nathan Kutz. Deep learning of conjugate mappings. *Physica D: Nonlinear Phenomena*, 427:133008, 2021.

[5] Jason J Bramburger and Giovanni Fantuzzi. Auxiliary functions as koopman observables: Data-driven analysis of dynamical systems via polynomial optimization. *Journal of Nonlinear Science*, 34(1):8, 2024.

[6] Jason J Bramburger and J Nathan Kutz. Poincaré maps for multiscale physics discovery and nonlinear floquet theory. *Physica D: Nonlinear Phenomena*, 408:132479, 2020.

[7] Shengze Cai, Zhiping Mao, Zhicheng Wang, Minglang Yin, and George Em Karniadakis. Physics-informed neural networks (PINNs) for fluid mechanics: A review. *Acta Mechanica Sinica*, pages 1–12, 2022.

[8] Elena Celledoni, Matthias J Ehrhardt, Christian Etmann, Robert I McLachlan, Brynjulf Owren, C-B Schonlieb, and Ferdia Sherry. Structure-preserving deep learning. *European journal of applied mathematics*, 32(5):888–936, 2021.

[9] Kathleen Champion, Bethany Lusch, J Nathan Kutz, and Steven L Brunton. Data-driven discovery of coordinates and governing equations. *Proceedings of the National Academy of Sciences*, 116(45):22445–22451, 2019.

[10] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.

[11] Salvatore Cuomo, Vincenzo Schiano Di Cola, Fabio Giampaolo, Gianluigi Rozza, Maizar Raissi, and Francesco Piccialli. Scientific machine learning through physics-informed neural networks: Where we are and what's next. *arXiv preprint arXiv:2201.05624*, 2022.

[12] Sigal Gottlieb, Zachary J Grant, Jingwei Hu, and Ruiwen Shu. High order strong stability preserving multiderivative implicit and imex runge–kutta methods with asymptotic preserving properties. *SIAM Journal on Numerical Analysis*, 60(1):423–449, 2022.

[13] Sigal Gottlieb, David Ketcheson, and Chi-Wang Shu. *Strong stability preserving Runge-Kutta and multistep time discretizations*. World Scientific, 2011.

[14] Ernst Hairer, Christian Lubich, and Gerhard Wanner. *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*, volume 31. Springer Science & Business Media, 2006.

[15] Jingwei Hu and Ruiwen Shu. A second-order asymptotic-preserving and positivity-preserving exponential runge–kutta method for a class of stiff kinetic equations. *Multiscale Modeling & Simulation*, 17(4):1123–1146, 2019.

[16] Alec J Linot and Michael D Graham. Deep learning to discover and predict dynamics on an inertial manifold. *Physical Review E*, 101(6):062209, 2020.

[17] Yuying Liu, J Nathan Kutz, and Steven L Brunton. Hierarchical deep learning of multiscale differential equation time-steppers. *Philosophical Transactions of the Royal Society A*, 380(2229):20210200, 2022.

[18] Yuying Liu, Colin Ponce, Steven L Brunton, and J Nathan Kutz. Multiresolution convolutional autoencoders. *Journal of Computational Physics*, 474:111801, 2023.

[19] Bethany Lusch, J Nathan Kutz, and Steven L Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature communications*, 9(1):4950, 2018.

[20] Eric J Parish and Kevin T Carlberg. Time-series machine-learning error models for approximate solutions to parameterized dynamical systems. *Computer Methods in Applied Mechanics and Engineering*, 365:112990, 2020.

[21] Tong Qin, Kailiang Wu, and Dongbin Xiu. Data driven governing equations approximation using deep neural networks. *Journal of Computational Physics*, 395:620–635, 2019.

[22] Francesco Regazzoni, Luca Dede, and Alfio Quarteroni. Machine learning for fast and reliable solution of time-dependent differential equations. *Journal of Computational physics*, 397:108852, 2019.

[23] Andrew Stuart and Anthony R Humphries. *Dynamical systems and numerical analysis*, volume 2. Cambridge University Press, 1998.

[24] Andy TS Wan, Alexander Bihlo, and Jean-Christophe Nave. Conservative methods for dynamical systems. *SIAM Journal on Numerical Analysis*, 55(5):2255–2285, 2017.

[25] Andy TS Wan and Jean-Christophe Nave. On the arbitrarily long-term stability of conservative methods. *SIAM Journal on Numerical Analysis*, 56(5):2751–2775, 2018.