AI4DifferentialEquations in Science

Workshop Summary

Over the past decade, the integration of Artificial Intelligence (AI) for scientific exploration has grown as a transformative force, propelling research into new realms of discovery. The ***AI4DifferentialEquations in Science*** workshop at ICLR 2024 invites participants on a dynamic journey at the interface of machine learning and computational sciences known as Scientific Machine Learning (SciML).

This workshop aims to unleash innovative approaches that harness the power of AI algorithms combined with computational mathematics to advance scientific discovery and problem solving. This enables us to push the boundaries of scientific computing beyond its traditional limits. Our goal is to delve into the latest AI advancements, particularly those that significantly enhance the efficiency of solving ordinary and partial differential equations (PDEs). These methods result in significant performance gains, which allow for solutions at high resolution that were previously unfeasible or required large amounts of computation. The ***AI4DifferentialEquations in Science*** workshop aims to unlock the full potential of data-driven approaches in advancing scientific frontiers in earth sciences, climate and computational fluid dynamics to name a few.

Key topics include but are not limited to:

- Exploration of novel applications of deep learning techniques in scientific simulations of partial or ordinary differential equations.
- Forward and inverse problems in PDEs to equation discovery, design optimization, and beyond, to witness the diverse applications of AI in scientific pursuits.
- Explainability and interpretability of AI models in scientific contexts.

The ***AI4DifferentialEquations in Science*** workshop will be an interactive event for researchers and practitioners at various stages of their careers with diverse backgrounds to come together and share their perspectives and experiences on leveraging AI techniques for their particular science problems. We aim to better bridge the gaps between the different fields of numerical analysis, machine learning and the sciences by sharing insights, methodologies, and challenges in harnessing AI's power for the greater good of scientific exploration.

Modality

We will offer a hybrid workshop format to accommodate both in-person and online attendance, thereby maximizing the outreach and impact of our workshop. We strongly encourage the invited speakers to attend in person if feasible for a more interactive experience with the workshop audience. To allow participants from various locations to engage in real-time, we will provide live-streamed sessions, and interactive discussions at the end of the invited talks and during the panel discussion. We will use **Zoom** for virtual talks and **Gather Town** for the hybrid poster sessions.

Tentative schedule

We aim to highlight recent and developing work as well as junior researchers, which is reflected in our invited speakers list as well as the number of contributed talks. We plan to include 6 invited talks (25 minutes + 5 minutes questions) and 4 contributed talks by authors of the submitted works. We aim to maximize interaction between participants, which we plan to accomplish via two poster sessions, the luncheon and the coffee breaks. To further foster interdisciplinary collaborations and discussions, the workshop will conclude with a panel including all of our invited speakers and other distinguished panelists. Topics of the panel discussions include the limitations of AI in tackling practical problems in science and best mechanisms to better facilitate meaningful interdisciplinary research among computational mathematicians, machine learning practitioners and scientists in application domains of earth science, climate and fluid dynamics. We will collect questions via an online form from the audience and online (streaming) viewers for our hybrid workshop to engage the online audience as well. Lunch and coffee breaks of the workshop will be coordinated with the official schedule of the conference. Our tentative schedule is as follows:

09:00	Introduction and Opening Remarks
09:15	Invited Talk 1
09:45	Contributed Talk 1
10:00	Invited Talk 2
10:30	Coffee Break
10:45	Poster Session 1
11:15	Invited Talk 2
11:45	Contributed Talk 2
12:00	Invited Talk 3
12:30	Lunch
13:30	Invited Talk 4
14:00	Poster Session 2
14:30	Contributed Talk 3
14:45	Coffee Break

15:00	Invited Talk 5
15:30	Contributed Talk 4
15:45	Invited Talk 6
16:15	Panel discussion
17:00	Coffee Break
17:15	Invited Talk 7
17:45	End

Invited Speakers/Panelists

Johannes Brandstetter (speaker, confirmed) is an Assistant Professor at Johannes Kepler University in Linz and a Guest Researcher at Microsoft Research Amsterdam, working on the intersection of deep learning and physics, aiming to accelerate partial differential equation solving and computational simulations. His research has spanned a range of topics from geometric deep learning, graph-structured learning, reinforcement learning, few-shot learning to particle and high-energy physics. Before starting his professorship, he was senior researcher at Microsoft, ELLIS postdoctoral researcher at the Amsterdam Machine Learning Lab (AMLAB) and at the Johannes Kepler University in Linz. He did his PhD at the CMS experiment at CERN investigating Higgs boson decays.

Steve Brunton (speaker, tentative) is a James B. Morrison Endowed Career Development Professor in Mechanical Engineering and Adjunct Professor of Applied Mathematics at the University of Washington. He is also a Data Science Fellow at the eScience Institute. He received his Ph.D. in Mechanical and Aerospace Engineering from Princeton University in 2012 and his B.S. in Mathematics with a minor in Control and Dynamical Systems from the California Institute of Technology in 2006. His research focuses on combining techniques in dimensionality reduction, sparse sensing, and machine learning for the data-driven discovery and control of complex dynamical systems. He is also interested in how low-rank coherent patterns that underlie high-dimensional data facilitate sparse measurements and optimal sensor and actuator placement for control. He is developing adaptive controllers in an equation-free context using machine learning. Specific applications in fluid dynamics include closed-loop turbulence control for mixing enhancement, bio-locomotion, and renewable energy. Other applications include neuroscience, medical data analysis, networked dynamical systems, and optical systems.

Youngsoo Choi (speaker, confirmed) is a computational math scientist in the Center of Applied Scientific Computing (CASC) under the Computing directorate at Lawrence Livermore National Laboratory (LLNL). His research focuses on developing efficient reduced order models for various physical simulations for time-sensitive decision-making multi-query problems, such as inverse problems, design optimization, and uncertainty quantification. His expertise includes various scientific computing disciplines. Together with his collaborators, he has developed various powerful model order reduction techniques, such as machine learning-based nonlinear manifold and space-time reduced order models for nonlinear dynamical systems. He has also developed the component-wise reduced order model optimization algorithm, which enables fast and accurate computational modeling tool for lattice-structure design. He is currently leading data-driven surrogate modeling development team for various physical simulations, with whom he developed the open source codes, libROM and LaghosROM. He is also involved with quantum computing research. He has earned his undergraduate degree for Civil and Environmental Engineering from Cornell University and his PhD degree for Computational and Mathematical Engineering from Stanford University. He was a postdoc at Sandia National Laboratories and Stanford University prior to joining LLNL in 2017.

Meire Fortunato (speaker, confirmed) is a research scientist at DeepMind, which she joined in 2016. Meire received her undergraduate and master's degrees in Mathematics from University of Campinas in Brazil (in 2008 and 2010, respectively); and her PhD in Mathematics in 2016 from the University of California, Berkeley. Meire is co-founder and a member of the general committee of Khipu.ai, an effort to support and strengthen AI research in Latin America. Meire has worked on a variety of machine learning topics, such as sequence modeling and exploration and understanding the role of memory in reinforcement learning agents. In the last two years, Meire's research focus has been on how to use machine learning for physics simulations -- with emphasis on graph neural networks architectures. This links to her mathematics background, as her PhD thesis was on the generation of high-order meshes by solving partial differential equations with finite element methods.

Paula Harder (speaker, confirmed) is a Ph.D. student in Computer Science at the Fraunhofer Institute for Industrial Mathematics (ITWM) and the Scientific Computing Group (SciComp) of the University of Kaiserslautern. She is supervised by Nicolas Gauger (SciComp) and Janis Keuper (ITWM). She also was an intern with David Rolnick at Mila, working on Physics-Constrained Climate Downscaling and a visiting researcher at the University of Oxford working on aerosol model emulation with Philip Stier and Duncan Watson-Parris. Her key research interest is Deep Learning applications for Climate and Weather Modeling. In addition, she works on robustness in Deep Learning on the topic of Adversarial Attacks and is interested in any applications of ML to climate, earth, and space science, as well as in the broad field of AI for Social Good. She is involved with the NASA/ESA Frontier Development Lab, where she worked as an ML Scientist and returned as a team lead in 2022, looking into aerosols from wildfires. She received a master's degree in Mathematics with a specialization in Numerical Analysis from the University of Tübingen, where she also got an education in Machine Learning. During her master's she did research projects at the German Climate Computation Center and a company for simulation of electrical networks in her spare time. After graduating she worked as a Development Engineer in the automotive industry, both on software development projects as well as data science applications.

Jakob Macke (speaker, confirmed) is a Professor for “Machine Learning in Science”. The W3 professorship has been set up as part of the Cluster of Excellence “Machine Learning: New Perspectives for the Sciences”. He is also an Adjunct Research Scientist at the Max Planck Institute for Intelligent Systems, Director of the Bernstein Center for Computational Neuroscience, and an ELLIS Fellow and member of the ELLIS Unit Tübingen.
Jakob studied mathematics at Oxford University, worked as a PhD student at the Max Planck Institute for Biological Cybernetics in Tübingen, as a postdoc at the Gatsby Unit at University College London, and as a Bernstein Fellow in Tübingen. He was a Max Planck Group Leader at the Caesar Research Centre in Bonn, a Professor at the Centre for Cognitive Science at TU Darmstadt, and from 2018 to 2020, Professor of Computational Neuroengineering at TU Munich. He was a member of the Young Academy at the German Academy of Sciences Leopoldina (2013-2018), and is currently a FENS Kavli Scholar of Excellence.

Christopher Rackauckas (speaker, confirmed) is Research Affiliate and Co-PI of the Julia Lab at the Massachusetts Institute of Technology, Director of Modeling and Simulation at Julia Computing and Creator and Lead Developer of JuliaSim, Director of Scientific Research at Pumas-AI and Creator / Lead Developer of Pumas and Lead Developer of the SciML Open Source Software Organization. His research and software is focused on Scientific Machine Learning (SciML): the integration of domain models with artificial intelligence techniques like machine learning. By utilizing the structured scientific (differential equation) models together with the unstructured data-driven models of machine learning, our simulators can be accelerated, our science can better approximate the true systems, all while enjoying the robustness and explainability of mechanistic dynamical models. Chris's recent work is focused on bringing personalized medicine to standard medical practice through the proliferation of software for scientific AI. Chris is at the cutting edge of mathematical methods for scientific simulation. He is the lead developer of the DifferentialEquations.jl solver suite along with over a hundred other Julia packages, earning him the inaugural Julia Community Prize, an outstanding paper award at the IEEE-HPEC conference on computational derivation for the efficient stochastic differential equation solvers, and front page features on many tech community sites. Chris' work on high performance differential equation solving is the engine accelerating many applications from the MIT-CalTech CLiMA climate modeling initiative to the SIAM Dynamical Systems award winning DynamicalSystems.jl toolbox (of which DifferentialEquations.jl was the runner-up). His work is credited for the 15,000x acceleration of NASA Launch Services simulations and recently demonstrated a 60x-570x acceleration over Modelica tools. For these achievements Chris received the United States Department of the Air Force Artificial Intelligence Accelerator Scientific Excellence Award.

Alex Townsend (speaker, confirmed) is an Associate Professor of Mathematics at Cornell University. He received his Ph.D. from the University of Oxford in 2014. His research interests include numerical analysis, scientific computing, and deep learning. He is interested in the study and development of numerical algorithms in applied mathematics. In particular, he works mainly in the following three areas: (1) Novel spectral methods for the solution of differential equations, (2) Low-rank techniques, and (3) Theoretical aspects of deep learning.

Kevin Carlberg (panelist, confirmed) is an AI Research Science Manager at Meta Reality Labs Research and an Affiliate Associate Professor of Applied Mathematics and Mechanical Engineering at the University of Washington. He leads a multidisciplinary research team focused on enabling the future of interaction in augmented and virtual reality by developing AI breakthroughs across the domains of embodied AI, multimodal machine learning, conversational AI, and AI-accelerated computational physics. His individual research combines concepts from machine learning, computational physics, and high-performance computing to drastically reduce the cost of simulating nonlinear dynamical systems at extreme scale. Previously, Kevin was a Distinguished Member of Technical Staff at Sandia National Laboratories in Livermore, California, where he led a research group of PhD students, postdocs, and technical staff in applying these techniques to a range of national-security applications in mechanical and aerospace engineering.

Marta D'Elia (panelist, confirmed) is a Principal Research and Computational Scientist at Pasteur Labs and an Adjunct Professor at the Institute of Computational and Mathematical Engineering at Stanford University. Previously, she was a Research Scientist at Meta and a Principal Member of the Technical Staff at Sandia National Laboratories, California. She works on the design and analysis of models and data-driven algorithms for the simulation of complex, multiscale and multiphysics problems. Her background and training have foundations in Numerical Analysis, Scientific Computing, Inverse Problems, Control and Optimization, and Uncertainty Quantification. In the past five years, she has focused on Scientific Machine Learning (SciML) and Deep Learning. She is an expert in Nonlocal/Fractional Modeling and Simulation with application to Continuum Mechanics, Subsurface Transport, Image Processing, and Turbulence. She has a master's degree in Mathematical Engineering from Politecnico di Milano (2007) and a PhD in Applied Mathematics from Emory University (2011).

Shirley Ho (panelist, confirmed) joined the Physics Department as Research Professor and as an Affiliated Faculty at Center for Data Science at NYU in 2021. She joined Simons Foundation in 2018 as leader of the Cosmology X Data Science group at CCA and in 2021, she assumed the role of CCA's interim director. Her research interests have ranged from fundamental cosmological measurements to exoplanet statistics to using machine learning to estimate how much dark matter is in the universe. Ho has broad expertise in theory, observation and data science. Ho's recent interest has been on understanding and developing novel tools in statistics and machine learning techniques, and applying them to astrophysical challenges. Her goal is to understand the universe's beginning, evolution and its ultimate fate. She earned her Ph.D. in astrophysical sciences from Princeton University in 2008 and her bachelor's degrees in computer science and physics from the University of California, Berkeley in 2004. She was a Chamberlain fellow and a Seaborg fellow at Lawrence Berkeley National Laboratory before joining Carnegie Mellon University in 2011 as an assistant professor. She became the Cooper Siegel Career Development Chair Professor and was appointed associate professor with tenure in 2016. She moved to Lawrence Berkeley Lab as a Senior Scientist in 2016. Since 2011, she has been a primary mentor to more than 35 postdoctoral fellows, 10 graduate students and 20 undergraduates in the fields of astrophysics, computer science and statistics. She has received several awards including NASA Group Achievement Award, Macronix Prize and Carnegie Science Award. She is also elected a Fellow by the International Astrostatistics Association.

Max Welling (panelist, confirmed) is a research chair in Machine Learning at the University of Amsterdam and a Distinguished Scientist at MSR. He is a fellow at the Canadian Institute for Advanced Research (CIFAR) and the European Lab for Learning and Intelligent Systems (ELLIS) where he also serves on the founding board. His previous appointments include VP at Qualcomm Technologies, professor at UC Irvine, postdoc at U. Toronto and UCL under supervision of prof. Geoffrey Hinton, and postdoc at Caltech under supervision of prof. Pietro Perona. He finished his PhD in theoretical high energy physics under supervision of Nobel laureate prof. Gerard‘t Hooft. He has served as associate editor in chief of IEEE TPAMI from 2011-2015, he serves on the advisory board of the Neurips foundation since 2015 and has been program chair and general chair of Neurips in 2013 and 2014 respectively. He was also program chair of AISTATS in 2009 and ECCV in 2016 and general chair of MIDL 2018. Max Welling is the recipient of the ECCV Koenderink Prize in 2010 and the ICML Test of Time award in 2021. He directs the Amsterdam Machine Learning Lab (AMLAB) and co-directs the Qualcomm-UvA deep learning lab (QUVA) and the Bosch-UvA Deep Learning lab (DELTA).

Organizers and Biographies

Shima Alizadeh is an Applied Scientist at AWS AI Labs at Amazon Web Services working on diverse machine learning problems including SciML/AI4Science, Multi-Objective Optimization (MOO), higher order optimization methods, personalization, and recommendations. Prior to joining AWS AI Labs she worked at Adobe as Machine Learning Engineer/Researcher for time-series forecasting problems with applications for intelligent marketing. She received her PhD in Mechanical Engineering from Stanford University with focus on high performance computing (HPC) for computational engineering, where she designed and implemented robust multi-scale models to study physical problems that involve transport phenomena and fluid dynamics. Shima’s passion lies at the interface of HPC for computational physics and Artificial Intelligence to solve problems that are governed by Physics based PDEs and ODEs with various applications such as climate research, transport phenomena, and engineering design optimization. At AWS AI Labs, by leveraging her expertise in HPC, numerical methods, machine learning, and optimization, she has been working on different problems such as Physics-Guided Neural Operators, Uncertainty Quantification (UQ) for SciML models, and ML models that respect physics constraints.

Neil Ashton is a Principal Computational Engineering Specialist at Amazon Web Services (AWS), as well as being a fellow of the Institute of Mechanical Engineering (IMechE). He leads the Amazon-wide Computational Fluid Dynamics (CFD) working group and acts as a CFD subject matter expert across the diverse departments using CFD at Amazon. Prior to joining AWS he obtained a Masters and Ph.D. at the University of Manchester focusing on turbulence modeling, and previous positions were at the University of Oxford, NASA Ames Research Center, and Lotus Formula 1 Team, as well as consultancy roles for a range of automotive and aerospace companies . His expertise is on applying and developing the next generation of CFD approaches with a particular focus on interplay between high-fidelity turbulence models, machine learning and high-performance computing. He created the Automotive CFD Prediction Workshop (whose 4th edition is happening in September 2024), and leads the hybrid RANS-LES technical working group for the 5th AIAA High-Lift Prediction CFD Workshop. In the past year his particular focus has been spearheading the development of industrialized versions of ML methods (e.g MeshGraphNet) that can be used to accelerate the engineering design process.

Vitus Benson is an ELLIS PhD student at Max Planck Institute for Biogeochemistry, ELLIS Unit Jena and ETH Zürich. He leverages machine learning to deepen our understanding of the Earth and climate sciences. In particular he is working on data-driven modeling of the carbon cycle. For this he is emulating an atmospheric transport model with deep learning. Furthermore he is interested in inverse modeling with PDEs. In addition he is a leading member of the EarthNet team, where he develops deep neural networks to predict the influence of weather on vegetation, his models can be used to estimate the localized impacts of weather extremes, towards improved early warning and to enhance humanitarian anticipatory action. Prior to his current work, he received a M.Sc. in Mathematical Physics from Leipzig University and a B.Sc. in Mathematics from the University of Göttingen. His research interests are broadly speaking Deep Learning, Biogeochemistry, Climate Risk Mitigation and Scientific Computing.

Boran Han is an Applied Scientist in the DeepEarth Group at AWS AI Research and Education. She earned her Ph.D. in biophysics from Harvard University under the guidance of Prof. Xiaowei Zhuang. Throughout her doctoral journey, she delved into pioneering super-resolution imaging techniques, with her research being published in esteemed journals such as Science and Cell, among others. Upon completion of her Ph.D., she joined Shell, where she leveraged deep learning to expedite seismic reconstruction, a complex inverse problem. Since joining AWS, Boran has redirected her research focus towards the development of multimodal foundation models specifically designed for scientific applications. One of her primary objectives has been to enhance the accuracy, robustness and explainability/interpretability of these foundation models. Additionally, she is dedicated to developing new foundation models tailored for scientific applications. Her commitment to merging advanced AI techniques with scientific challenges places her at the forefront of this dynamic interdisciplinary field.

George Karypis is a Distinguished McKnight University Professor in the Department of Computer Science & Engineering at the University of Minnesota. He received his B.S. and Ph.D. in Computer Science from the University of Minnesota. He joined the Department of Computer Science & Engineering in 1999. He was promoted to full professor in 2009. He is the co-author of two books and has contributed to the development of multiple software libraries. Additionally, he serves on the program committees of many conferences and workshops, and on the editorial boards of the IEEE Transactions on Big Data, ACM Transactions on Knowledge Discovery from Data, Data Mining and Knowledge Discovery, Journal of Data Science and Analytics, Social Network Analysis and Data Mining Journal, International Journal of Data Mining and Bioinformatics, the journal on Current Proteomics, Advances in Bioinformatics, and Biomedicine and Biotechnology. He is an IEEE Fellow.

Aditi Krishnapriyan is an Assistant Professor of Chemical and Biomolecular Engineering; Electrical Engineering and Computer Sciences at the University of California, Berkeley. She is a Bruce and Susan Stangeland Processor and Member of the Berkeley AI Research (BAIR). Her research interests are focused on developing machine learning methods that are motivated by the opportunities and challenges in science and engineering. Some of the areas of exploration include approaches to incorporate physical inductive biases into ML models to improve generalization, the advantages that ML can bring to classical physics-based numerical solvers (such as through end-to-end differentiable frameworks and implicit layers), and better learning strategies for distribution shifts in the physical sciences. Her foundational research is informed by and grounded in applications in physics, fluid and molecular dynamics, materials design, climate science, and other related areas. Her work also includes interfacing with other fields including numerical methods, dynamical systems theory, quantum mechanical simulations, computational geometry, and optimization.

Danielle Maddix Robinson is a Senior Applied Scientist in the Machine Learning Forecasting Group within AI Research and Education at Amazon Web Services (AWS). She graduated with her PhD in Computational and Mathematical Engineering from the Institute of Computational and Mathematical Engineering (ICME) at Stanford University. She was advised by Professor Margot Gerritsen and developed robust numerical methods to remove spurious temporal oscillations in the degenerate Generalized Porous Medium Equation. She is passionate about the underlying numerical analysis, linear algebra and optimization methods behind numerical PDEs and applying these techniques to deep learning. She joined AWS in 2018 shortly after graduating, and has been working on developing statistical and deep learning models for time series forecasting. In this past year, she has been leading the research initiative on developing models for physics-constrained machine learning for scientific computing on the DeepEarth team. In particular, she has researched how to apply ideas from numerical methods, e.g., finite volume schemes, to improve the accuracy of black-box ML models for ODEs and PDEs with applications to epidemiology, aerodynamics, ocean and climate models.

Michael W. Mahoney is at the University of California at Berkeley in the Department of Statistics and at the International Computer Science Institute (ICSI). He is also a faculty scientist at the Lawrence Berkeley National Laboratory. He works on algorithmic and statistical aspects of modern large-scale data analysis. Much of his recent research has focused on large-scale machine learning, including randomized matrix algorithms and randomized numerical linear algebra, geometric network analysis tools for structure extraction in large informatics graphs, scalable implicit regularization methods, computational methods for neural network analysis, physics informed machine learning, and applications in genetics, astronomy, medical imaging, social network analysis, and internet data analysis. He received his PhD from Yale University with a dissertation in computational statistical mechanics, and he has worked and taught at Yale University in the mathematics department, at Yahoo Research, and at Stanford University in the mathematics department. Among other things, he is on the national advisory committee of the Statistical and Applied Mathematical Sciences Institute (SAMSI), he was on the National Research Council's Committee on the Analysis of Massive Data, he co-organized the Simons Institute's fall 2013 and 2018 programs on the foundations of data science, he ran the Park City Mathematics Institute's 2016 PCMI Summer Session on The Mathematics of Data, and he runs the biennial MMDS Workshops on Algorithms for Modern Massive Data Sets. He is the Director of the NSF/TRIPODS-funded FODA (Foundations of Data Analysis) Institute at UC Berkeley.

Markus Reichstein is Director of the Biogeochemical Integration Department at the Max-Planck-Institute for Biogeochemistry. His main research interests revolve around the response and feedback of ecosystems (vegetation and soils) to climatic variability with an Earth system perspective. Of specific interest is the interplay of climate extremes with ecosystem and societal resilience. He is addressing these topics with a combination of artificial intelligence and system modelling approaches to exploit the wealth of experimental, ground- and satellite-based Earth observations. Markus studied Ecology with Botany, Chemistry and Computer Science as minor and obtained his PhD in Plant Ecology at the University of Bayreuth, Germany. Since 2013, Markus has been Professor for Global Geoecology at the FSU Jena, and founding Director at the Michael-Stifel-Center Jena for Data-driven and Simulation Science. He is fellow within the excellence network ELLIS (European Laboratory for Learning and Intelligent Systems) and founding director of the ELLIS Unit Jena established in 2021. He has been serving as lead author of the IPCC special report on Climate Extremes (SREX), as member of the German Committee Future Earth on Sustainability Research, and the Thuringian Panel on Climate. Recent awards include the Piers J. Sellers Mid-Career Award by the American Geophysical Union (2018), an ERC Synergy Grant (2019) and the Gottfried Wilhelm Leibniz Preis (2020).

Yuyang (Bernie) Wang is a Principal Machine Learning Scientist in Amazon AI Labs, working mainly on large-scale probabilistic machine learning with its application in Forecasting. He received his PhD in Computer Science from Tufts University, MA, US and he holds an MS from the Department of Computer Science at Tsinghua University, Beijing, China. His research interests span statistical machine learning, numerical linear algebra, and random matrix theory. In forecasting, Yuyang has worked on all aspects ranging from practical applications to theoretical foundations.

Andrew Gordon Wilson is an Associate Professor in the Courant Institute of Mathematical Sciences and Center for Data Science at New York University. His work is focused on developing a prescriptive understanding of model construction and generalization, often working with probabilistic and Bayesian methods. His research agenda includes uncertainty representation, scalable inference, physics-inspired machine learning, neural networks for differential equations, and many scientific applications, including protein engineering, spatiotemporal modelling, general relativity, and medical diagnosis. He has organized over 10 highly successful workshops across NeurIPS and ICML, including a major NeurIPS 2017 symposium on interpretable machine learning. Andrew also co-organized the NeurIPS 2021 competition for approximate inference in Bayesian deep learning, has been the EXPO Chair for ICML 2020, 2021, and is the Tutorial Chair for NeurIPS 2022, 2023, and will be the Workshop Chair for ICML 2024. He was also the ICML 2020 tutorial speaker in Bayesian deep learning.

Mahkameh (Maggie) Zarekarizi is a Research Scientist in the World Wide Sustainability group at Amazon Web Services (AWS). She is a climate science subject matter expert across business units that work on or are impacted by environmental hazards. Her research involves the projection and uncertainty quantification of climate-related risks across various hazards and geographies. She uses Decision Making Under Deep Uncertainty (DMDU) techniques to help stakeholders make robust decisions to manage and mitigate their climate risk. Before joining Amazon, Maggie was a data scientist at Jupiter Intelligence, where she led the development of high-resolution global flood and drought risk models leveraging data-driven and machine learning approaches. She completed a postdoctoral scholarship at the Earth and Environmental Systems Institute at Penn State and earned her PhD from the Maseeh College of Engineering at Portland State University. Her PhD research involved assimilating real-time satellite observations into Earth system models to improve their accuracy in forecasting environmental hazards.

Anticipated audience size

We expect to have approximately 100 in person attendees.

Plan to get an audience for a workshop

We will advertise our workshop on our mailing lists at our different institutions. We will also advertise on our Twitters and LinkedIns and MLNews (ml-news@googlegroups.com) to maximize the reach-out to the community.

Diversity commitment

Our commitment to diversity extends across various dimensions, including but not limited to ethnicity, age, gender, background, scientific seniority, and professional expertise. Our organizing team contains professionals with a wide range of backgrounds and perspectives, ensuring a rich and inclusive environment. We deliberately selected a speaker lineup that reflects a diverse experience with different scales of scientific seniority including PhD candidates, senior researchers from industry and national laboratories, as well as assistant and full professors.

Access

We will publish the accepted papers and talk abstracts before the event and the slides of the speakers after the event on the workshop website. We will include a bibliography of the most relevant research papers to facilitate interdisciplinary research of sharing ideas between the fields of numerical analysis, scientific computing and the sciences. Finally, we aim to record the workshop and publish it online assuming that we can obtain necessary funds.

Funding

Upon acceptance, we will solicit sponsorships from companies such as Amazon, DeepMind, and Microsoft. The funds will be used for accessibility (record and stream the workshop talks), awards (best poster and talk), and to create student travel grants.

Previous related workshops

This is the first edition of the workshop AI4DifferentialEquations in Science. This workshop focuses on utilizing ML methods for computational sciences with applications in climate research, Earth science and computational fluid dynamics. Previous workshops on related topics include: Physics4ML (ICLR 2019), Machine Learning for Physical Sciences (NeurIPS

20202-2023), AI4Science (NeurIPS 2023), AI for Scientific Discovery: From Theory to Practice (NeurIPS 2023), Tackling Climate Change with Machine Learning (NeurIPS 2023), Symbiosis of Deep Learning and Differential Equations (NeurIPS 2022). In contrast to prior workshops, this workshop is uniquely positioned to bring together researchers that study varying methodologies and perspectives on methods in numerical analysis, scientific computing, and machine learning and how best to apply them to various scientific disciplines. There will be an emphasis on the computational mathematical tools powering these methods. We target both scientists who aim to use ML methods and ML practitioners who are interested in harnessing ML's power for scientific applications. We aim to maximize diversity and bring together researchers from various fields to better foster and enable effective interdisciplinary applications of machine learning and computational mathematics in the sciences.