AI&PDE: ICLR 2026 WORKSHOP ON AI AND PARTIAL DIFFERENTIAL EQUATIONS

ABSTRACT

Partial Differential Equations (PDEs) are foundational to modeling complex phenomena across the natural sciences and engineering, from fluid dynamics and quantum systems to climate modeling and materials science. Despite their ubiquity, solving PDEs remains computationally intensive, especially in high-dimensional, multi-physics, and uncertain regimes. Recent advances in machine learning—such as neural operators, physics-informed networks, and foundation models—offer transformative potential to accelerate and generalize PDE solutions. However, realizing this promise requires addressing critical challenges in representation, stability, generalization, and benchmarking.

The AI&PDE-ICLR-2026 workshop will convene researchers from machine learning, applied mathematics, physics, and engineering to explore the future of AI-driven PDE modeling. We aim to (1) define the roadmap toward foundation models for PDEs, (2) investigate next-generation representations and architectures, and (3) foster a globally inclusive community. The program will feature invited talks, contributed papers, and themed tracks, including a full papers track for mature research and a tiny papers track for emerging ideas. By bridging disciplines and promoting open benchmarks and datasets, AI&PDE-ICLR-2026 will catalyze progress toward scalable, general-purpose AI solvers for PDEs.

1 INTRODUCTION

Partial Differential Equations (PDEs) are the universal language of the natural sciences Evans (2022). They describe phenomena as diverse as turbulent flows in fluid mechanics Sagaut & Cambon (2008), electronic transport in materials Lundstrom (2002), climate dynamics Palmer (2022), plasma physics Manfredi et al. (2021), and quantum many-body systems Bruus & Flensberg (2004); Pethick & Smith (2008). Yet, despite their centrality, solving PDEs remains one of the most resource-intensive tasks in computational science Karniadakis et al. (2021); Meng et al. (2025). Traditional numerical methods, including finite difference, finite element, spectral, and multigrid solvers, have advanced considerably; however, they often scale poorly with dimensionality, complex boundary conditions, and multiphysics couplings Chen et al. (2022); Galaris et al. (2022). As scientific challenges grow in ambition, from exascale simulations of fusion plasmas to uncertainty quantification in climate models, new computational paradigms are urgently needed Fasoli et al. (2016).

Recent breakthroughs in machine learning offer a transformative opportunity. Neural PDE solvers, physics-informed learning frameworks Karniadakis et al. (2021), and operator-based foundation models Kovachki et al. (2023) are emerging as powerful tools that can augment or even replace classical solvers. Approaches such as Physics-Informed Neural Networks (PINNs) Cai et al. (2021), Deep Operator Networks (DeepONets) Wang et al. (2022), Fourier Neural Operators (FNOs) Kovachki et al. (2021), and sequence-based models (e.g., Mamba-inspired architectures) Cheng et al. (2024); Gao et al. (2025) have shown promise in learning solution operators, generalizing across boundary conditions, and accelerating simulations by orders of magnitude. Moreover, foundation-model approaches trained on vast PDE datasets are beginning to demonstrate cross-domain generalization, suggesting a path toward reusable, pre-trained scientific models analogous to those that revolutionized natural language processing and vision Herde et al. (2024); Wiesner et al. (2025); Shen et al. (2024); Sun et al. (2025).

At the same time, AI for PDEs faces unique challenges: ensuring stability and physical consistency; integrating heterogeneous discretizations such as grids, meshes, graphs, and spectral bases; handling

multi-scale dynamics; and establishing benchmarks and evaluation standards that bridge communities of machine learning, numerical analysis, and domain sciences. These challenges demand a forum where researchers from diverse backgrounds can come together, share insights, and chart the future of AI-driven PDE modeling.

The proposed **AI&PDE-ICLR-2026** workshop will provide that forum. Building on the success of related workshops at NeurIPS and ICML, this will be the first ICLR venue dedicated to the intersection of PDEs and AI. Our goals are threefold:

1. To define the roadmap toward foundation models for PDEs.
2. To explore next-generation representations and architectures for PDE learning.
3. To foster a globally inclusive community spanning applied mathematics, physics, engineering, and machine learning.

By convening experts across disciplines, AI&PDE-ICLR-2026 will highlight both the remarkable progress already made and the open technical questions that must be addressed to fully realize AI's impact on PDE-driven science and engineering.

2 AI&PDE-ICLR-2026 Proposed Workshop

The **AI&PDE-ICLR-2026** workshop will provide a dedicated forum for researchers in machine learning, applied mathematics, physics, and engineering to address the challenges and opportunities of applying AI to Partial Differential Equations (PDEs).

Technical Goals & Problems to Solve

We frame the program around three central and urgent research challenges:

1. **How can we build a foundation model for PDEs?** While neural operators, PINNs, and hybrid solvers have shown promise, an open challenge remains: designing scalable, general-purpose foundation models that can transfer across PDE families, geometries, and boundary conditions. Addressing this problem requires cross-disciplinary expertise and shared benchmarks.
2. **What are next-generation representations of PDEs?** PDE systems require diverse representations (structured grids, irregular meshes, spectral expansions, graphs). The right choice of representation directly impacts stability, accuracy, and generalization. We seek advances in operator learning, multimodal embeddings, mesh-free methods, and geometric deep learning to overcome these bottlenecks.
3. **How do we build shared datasets and benchmarks for PDE learning?** Unlike fields such as vision and NLP, PDE datasets are fragmented across domains (fluids, quantum systems, elasticity, climate) with heterogeneous discretizations and varying levels of physical fidelity. Standardized benchmarks that include realistic boundary conditions, multi-physics couplings, and uncertainty quantification are largely missing. A key challenge is curating open, diverse, and multimodal datasets that can serve as the foundation for reproducible research and the evaluation of PDE foundation models.

By explicitly highlighting these problems, the workshop encourages participants to propose novel methods, new datasets, benchmarks, and theoretical insights that push the frontier of PDE learning.

We commit to a program that **centers contributed work**, with oral spotlights, poster sessions, and interactive discussion time. Roughly half of the program will be devoted to invited talks and panels, and the other half to contributed papers and posters. To enhance visibility, accepted works will be featured in spotlight sessions interleaved with invited keynotes.

Full Papers Track

We will host a **Full Papers Track** (up to 9 pages, excluding references) to encourage comprehensive contributions. This track will:

- Welcome mature research on neural PDE solvers, operator learning, physics-informed methods, datasets, and benchmarks.
- Allow detailed methodological, theoretical, and empirical studies with rigorous evaluation.
- Ensure visibility for these contributions through oral presentations, spotlights, and poster sessions.

This track provides a venue for substantial work that may not yet be ready for submission to the ICLR main conference but benefits from in-depth community feedback.

Tiny and Short Papers Track

We will host a **Tiny Papers Track** (2–4 pages) modeled on the ICLR initiative. This track will:

- Welcome late-breaking results, negative findings, benchmarks, and small-scale experiments.
- Provide opportunities for students and junior researchers outside the traditional ML conference circuit.
- Exclude AI-generated papers, in compliance with ICLR policy, while permitting AI-assisted writing with human oversight.

This track ensures accessibility and inclusivity, while prioritizing timely feedback from the ICLR community.

Submission Timeline

Date (AOE)	Milestone
8 Sep 2025	Workshop application opens
10 Oct 2025 (11:59pm)	Workshop application deadline
1 Dec 2025	Workshop acceptance notifications
8–12 Dec 2025	Finalize CFP text, website, OpenReview venue; recruit reviewers
15 Dec 2025	Call for Papers (CFP) announced; PC onboarding begins
6–10 Jan 2026	Reviewer assignment & bidding period
30 Jan 2026	Suggested submission deadline for workshop contributions
31 Jan–20 Feb 2026	Review period
21–28 Feb 2026	Final decisions and meta-reviews
1 Mar 2026 (11:59pm)	**Global author notification deadline (mandatory)**; all accepted papers public on OpenReview
2–10 Mar 2026	Camera-ready updates (minor fixes, links, artifacts)
11 Mar 2026 (11:59pm)	**Import workshop program & accepted papers to `iclr.cc`**

Table 1: End-to-end submission timeline aligned with ICLR 2026 workshop requirements. All times are AoE.

Community-Building Initiatives

- **Themed Track: Multi-Physics & Multimodal PDE Learning**: A focused track on PDEs that span domains (e.g., fluid-structure interaction, thermo-mechanics, and climate PDEs with multimodal data).
- **Mentorship**: Organizers will encourage informal mentoring during poster sessions and themed tracks, pairing early-career participants with senior researchers.
- **Global South Engagement:** The organizing committee includes members from Latin America and strong collaborations with institutions such as IMPA, IMPA Tech, UFRJ, CBPF, and Petrobras, ensuring active participation from underrepresented regions in scientific ML. These efforts aim to promote equitable access through hybrid participation, and community-led discussions.

- **Post-Workshop Focus Collection: IOP Machine Learning: Science and Technology (MLST)**: To extend the impact of the workshop, we have partnered with the journal *Machine Learning: Science and Technology* (IOP MLST) to organize a dedicated **Focus Collection** featuring selected contributions and invited papers arising from the AI&PDE community. This initiative will provide a formal venue for publishing extended versions of high-quality workshop papers, fostering continued collaboration among researchers from machine learning, applied mathematics, and the physical sciences.
- **Pre-Workshop Outreach Event (Brazil):** In collaboration with local partners, including the PUC–Behring Institute for Artificial Intelligence, IMPA, IMPA Tech, UFRJ, CBPF, and PETROBRAS, we will organize a pre-workshop event to engage students and early-career researchers. This event will include tutorials, lightning talks, and networking sessions designed to build local capacity and strengthen Brazil's role in the global AI&PDE research ecosystem.

Advertising and Outreach Plan for AI&PDE Workshop at ICLR 2026

The AI&PDE Workshop at ICLR 2026 aims to bring together researchers and professionals working at the intersection of Artificial Intelligence and Partial Differential Equations. To maximize visibility and engagement, we propose a multi-channel outreach strategy targeting academia, industry, and the broader AI community.

Workshop Website

A dedicated webpage will be created with the following sections:

- **Home Page**: Overview of the workshop, objectives, and relevance.
- **Speakers & Panelists**: Bios, photos, affiliations, and talk titles.
- **Schedule**: Agenda with session times and formats.
- **Call for Participation**: Submission guidelines and registration info.
- **Contact & Updates**: Newsletter signup and social media links.

Academic Outreach

- Faculty ambassadors to promote the workshop in classes and seminars. The use of workshop ambassadors at academic and research institutions involves appointing individuals responsible for promoting the workshop within their organizations and assisting with submissions. Confirmed institutions with ambassadors include: Caltech, CBPF, CENPES-Petrobras, Dallara, EPFL, ETH Zurich, FGV - RJ, FIOCRUZ, IBM, IMPA, IMPA Tech, Johns Hopkins University, Lancaster University, LNCC, MIT, Nvidia, PUC-Behring Institute for Artificial Intelligence, UCLA, UCSD, UFBA, UFF, UFRGS, UFRJ, Unicamp, University of Georgia, and USP.

Industry Engagement

- Internal newsletters via IBM, DALLARA, PETROBRAS, and NVIDIA.
- LinkedIn posts from company and personal profiles.
- Pre-event webinars featuring workshop speakers.

Social Media Campaign: Platforms

LinkedIn, X (Twitter), Instagram, and YouTube.

Cross-Promotion

- Submit content to ICLR's official channels.
- Announcements on partner institution websites.
- Blurbs in newsletters such as AIhub and ML Collective.

PREVIOUS RELATED WORKSHOPS

Recent years have seen growing interest in AI for scientific modeling across top venues. At **ICLR 2024**, the *AI4DifferentialEquations in Science* workshop brought together methods for ODEs/PDEs and highlighted open challenges in scalability and generalization, confirming strong community demand for this area.[1] Earlier, **ICLR 2023** hosted *Physics for Machine Learning*, underscoring links between representation learning and physical modeling, including neural PDE solvers.[2] At **NeurIPS 2022**, *The Symbiosis of Deep Learning and Differential Equations II* focused on neural DEs and PINNs.[3] **ICML 2023**'s *Synergy of Scientific and Machine Learning Modeling (SynS & ML)* further emphasized hybrid scientific–ML paradigms.[4]

Building upon these efforts, *AI&PDE–ICLR 2026* represents a natural evolution—yet with a distinct emphasis. It is the first ICLR workshop dedicated specifically to **foundation models for PDEs**, **cross-discretization operator learning**, and the **creation of shared multimodal benchmarks**. The workshop also pioneers the integration of a post-event *IOP MLST Focus Collection*, ensuring long-term visibility and scholarly continuity. Beyond its technical agenda, *AI&PDE–ICLR 2026* stands out for its commitment to **open science, reproducibility, and global inclusion**, fostering collaboration across disciplines and geographic regions.

COMPLIANCE WITH ICLR REQUIREMENTS

- **Notification Deadline**: All acceptance decisions (oral, poster, full papers, and tiny papers) will be communicated via OpenReview before 1 March 2026 (AOE).
- **Conflict of Interest Management**: Organizers will not act as reviewers for submissions from their own institutions. OpenReview COI tools will be leveraged, and a program committee with more than 30 institutions will ensure fairness.
- **In-Person Plan**: The workshop will be hosted on-site at ICLR 2026. We will record all talks and make posters and papers available online to maximize accessibility for participants who face visa or funding barriers.
- **LLM Usage Policy**: In line with ICLR 2026 policies, we explicitly prohibit AI-generated papers in tiny/short tracks. AI assistance is allowed but must be transparently acknowledged, with human authors retaining responsibility.
- **Audience Size & Engagement**: We anticipate approximately 150 participants, drawn from ML, physics, and applied math communities. Engagement will be fostered via structured Q&A, poster networking, and panel discussions.
- **Diversity Commitment**: The organizing team and invited speakers will represent diverse perspectives across gender, geography, discipline, and seniority.

ANTICIPATED AUDIENCE & PARTICIPATION

We anticipate an audience of approximately 150 participants from academia, research labs, and industry, spanning machine learning, applied mathematics, physics, and engineering. We expect a balanced mix of early-career researchers and senior experts. Engagement will be fostered through open Q&A sessions, interactive poster discussions, and mentorship meetups during lunch and coffee breaks. We will also encourage the submission of "tiny papers" from students and institutions under-resourced to maximize accessibility and inclusion.

VIRTUAL ACCESS & OUTCOMES

All talks, panels, and poster sessions will be recorded and made available on the workshop website shortly after the event. Accepted papers and posters will be archived on OpenReview, ensuring long-term accessibility. In addition, we will release slides, benchmark datasets, and open-source code

[1] https://ai4diffeqtnsinsci.github.io/
[2] https://physics4ml.github.io/
[3] https://dlde-2022.github.io/
[4] https://syns-ml.github.io/2023/

linked to workshop contributions when available. The post-workshop IOP MLST Focus Collection will further consolidate outcomes and sustain community collaboration.

LOGISTICS & CLARITY

Scope, Page Limits, and Tracks. We will follow ICLR 2026 workshop policies. The workshop will host two tracks: (i) **Full Papers** (up to **9 pages** excluding references; ICLR format), and (ii) **Tiny Papers** (2–4 pages; ICLR Tiny policy; no AI-generated papers, AI-assistance allowed with disclosure). Submissions must include a *limitations* paragraph and an ethics statement when applicable.

Review Process and COI. Reviews will be managed in **OpenReview**. Each paper receives **3 reviews** plus a **meta-review**. We will run a **bidding** phase, followed by **automatic and manual assignment** with topical matching. Conflicts are enforced via OpenReview's COI graph (institutional, advisor/advisee, coauthors within 3 years). Organizers will not handle submissions from their own institutions or collaborators and will not give talks at the workshop. **Desk-reject** criteria: policy violations (e.g., anonymity breaches), over-length, missing artifact for "benchmark" papers, or non-scientific content.

Calibration & Mentored Reviewing. Before reviews open, we will provide 2 calibration papers per area (operator learning; stability/conservation; discretization-robust models; benchmarking/UQ) with guidance rubrics. First-time reviewers can opt into a **mentored-review** path led by senior PC members.

Decision Policy and Visibility of Contributed Work. Decisions balance technical soundness, relevance to AI&PDE themes, clarity, and potential for discussion. Contributed work is centered: $\geq$**50%** of agenda time is devoted to contributed oral sessions, posters, and Q&A. All accepted papers receive a poster; top-ranked receive spotlights or short orals.

Artifacts, Reproducibility, and Benchmarks. For papers proposing datasets/benchmarks or reporting on PDE FMs, we **require** an artifact checklist (code or evaluation scripts; data or data-access instructions; license; compute footprint; seeds). We will host a public GitHub org and link DOIs where available. We will publish a minimal **evaluation kit** (stability/constraint violation metrics, OOD generalization splits, and UQ protocols) and maintain a **leaderboard** for seeded tasks.

Hybrid Access and Remote Participation. The workshop is in-person by default. We will **livestream via Zoom** with moderated remote Q&A (Slido backchannel). Poster authors may optionally upload a **3-minute lightning video** and a one-page PDF teaser; all talks are recorded and posted on the workshop website.

Inclusion and Accessibility. We will reserve **student-only lightning slots** in the Tiny track. We will provide captioned recordings and share slides in accessible formats. Visa/attendance guidance will be linked from the website.

Risk Mitigation. For unavoidable speaker conflicts, we will use pre-recorded talks with **live Q&A** or substitute confirmed alternates. If network issues arise, we will **mirror slides/videos** on the website and OpenReview and collect **asynchronous Q&A** for 7 days post-event.

Timeline Compliance. Our submission and decision schedule (Sec. 2) satisfies the **global notification deadline** of **1 March 2026 (AOE)**; the program is finalized and uploaded by **11 March 2026**.

3 TENTATIVE SCHEDULE

The AI&PDE-ICLR-2026 workshop will run for a full day during ICLR 2026. The agenda is designed to balance invited content with contributed work, while ensuring time for interactive panels, poster sessions, and community-building activities.

Time	Session
08:00–08:15	**Opening Remarks**: Workshop overview and goals
08:15–09:00	**Keynote 1: Anima Anandkumar (Caltech)**
09:00–09:30	Invited Talk: **Clécio R. Bom (CBPF)**
09:30–10:30	Contributed Oral Session I (Full Papers)
10:30–11:00	Coffee Break & Poster Session I (Full, Short, Tiny Papers)
11:00–11:30	Invited Talk: **Maximilian Herde (ETH Zürich)**
11:30–12:00	Invited Talk: **Jingmin Sun (Johns Hopkins/MIT)**
12:00–12:30	Contributed Oral Session II (Full Papers)
12:30–13:30	Lunch Break & Mentorship Meetups at Posters
13:30–14:15	**Keynote 2: Rose Yu (UC San Diego & Amazon Scholar)**
14:15–14:45	Invited Talk: **Hayden Schaeffer (UCLA)**
14:45–15:15	Contributed Oral Session III (Tiny & Short Papers)
15:15–15:45	Coffee Break & Poster Session II
15:45–16:30	**Panel: The Future of AI and PDEs: An Industrial View**
16:30–16:50	Demo & Reproducibility Showcase (Open-source tools, benchmarks, datasets)
16:50–17:35	**Keynote 3: Cristiano Malossi (IBM Research Zurich)**
17:35–17:45	**Closing Remarks**: next steps for the AI&PDE research community

Table 2: Tentative workshop agenda with two additional invited speakers, maintaining balance between keynotes, invited talks, contributed oral sessions, and community-building activities.

4 Invited Speakers

Our invited speakers represent a diverse set of perspectives across machine learning, applied mathematics, physics, and engineering. Below we include short bios to highlight their expertise and relevance to the AI&PDE community. Most of the invited speakers listed below have confirmed their participation. The lineup reflects diversity across geography, career stage, gender, and expertise, in line with ICLR workshop guidelines.

Anima Anandkumar – California Institute of Technology – (Website, Google Scholar) CONFIRMED
Anima Anandkumar is a Bren Professor at Caltech. She previously was a Senior Director of AI Research at NVIDIA and Principal Scientist at Amazon Web Services. She received her B.Tech from the Indian Institute of Technology Madras, and her Ph.D. from Cornell University. She did her postdoctoral research at MIT and an assistant professorship at the University of California Irvine. She has received several honors such as the IEEE Fellowship, Alfred. P. Sloan Fellowship, NSF CAREER Award, and Faculty Fellowships from Microsoft, Google, Facebook, and Adobe. She is also part of the World Economic Forum's Expert Network.

Clécio R. Bom – Centro Brasileiro de Pesquisas Físicas (CBPF) – (Website, Google Scholar) CONFIRMED
Clécio R. Bom is Professor and Technology Researcher at the Brazilian Center for Physics Research (CBPF), where he leads the Laboratory for Artificial Intelligence Applied to Physics (LAB-IA). His current work is focused on pushing the frontiers of AI in physics applications, including the inverse modelling with Neural Posterior Estimation, Physics-Informed Neural Networks and Agentic AI for physics in multiple fields including Astrophysics and Geophysics..

Cristiano Malossi – IBM Research – (Website, Google Scholar) CONFIRMED
Cristiano Malossi is a Principal Research Scientist and Manager of the Frontiers of Computing and Simulation group at the IBM Research Laboratory in Zurich. From 2025, Cristiano is leading research on AI for PDEs applied to Physics and Engineering. Previously, between 2020-2024, Cristiano has led IBM's global research and innovation strategy around Enterprise Visual Inspection, with a focus on inspection of large-scale infrastructures. His team designs, develops, and productizes scalable AI

cloud services for detection of small and rare defects in high-resolution data, bridging research with industrial deployment.

Hayden Schaeffer – University of California – (Website, Google Scholar) Tentative
Hayden Schaeffer is the Director of Applied Mathematics and a Professor of Mathematics at the University of California, Los Angeles. His research is in mathematical and scientific machine learning, differential equations, randomization, and modeling. He has received an NSF CAREER award and an AFOSR Young Investigator Award. Previously, he was an NSF Mathematical Sciences Postdoctoral Research Fellow, a von Karman Instructor at Caltech, a UC President's Postdoctoral Fellow at UC Irvine, an NDSEG Fellow, and a Collegium of University Teaching Fellow at UCLA.

Jingmin Sun – Johns Hopkins University – (Website, Google Scholar) CONFIRMED
Jingmin Sun is a postdoctoral fellow in the Department of Applied Mathematics and Statistics at Johns Hopkins University (JHU), working with Prof. Mauro Maggioni. She obtained her Ph.D. in Mathematical Science from Carnegie Mellon University (CMU), working with Prof. Hayden Schaeffer (UCLA), and a B.S. degree in Mathematical Science from Rensselaer Polytechnic Institute (RPI). Her research interests lie in Mathematical machine learning, differential equations, and optimization.

Maximilian Herde – ETH Zürich – (Website, Google Scholar) to be confirmed
Maximilian Herde is a doctoral researcher at the Computational and Applied Mathematics Laboratory (CAMLab) group, ETH Zürich, working under the supervision of Prof. Siddhartha Mishra. His research focuses on scientific machine learning, operator learning, and foundation models for partial differential equations. He is a coauthor of RIGNO and Poseidon, frameworks for robust and generalizable PDE operator learning. Maximilian received the ETH Medal in 2025 for his master's thesis"On Foundation Models for Partial Differential Equations".

Rose Yu – UC San Diego & Amazon Scholar – (Website, Google Scholar) CONFIRMED
Rose Yu is an Associate Professor at UC San Diego in the Department of Computer Science and Engineering and an Amazon Scholar. She is a primary faculty member with the AI Group. Her research interests focus on machine learning for large-scale spatiotemporal data, and she is particularly excited about AI for scientific discovery. She has received the Presidential Early Career Award for Scientists and Engineers (PECASE), the DARPA Young Faculty Award, the ECASE Award, the NSF CAREER Award, the Hellman Fellowship, and multiple faculty awards from Sony, JP Morgan, Meta, Google, Amazon, and Adobe. She was named one of MIT Technology Review's Innovators Under 35 in AI.

5 PANELISTS

In addition to invited speakers, we will host a panel discussion with confirmed participants from academia and industry. This panel brings together industry leaders to discuss how artificial intelligence is transforming the way complex PDE-driven problems are tackled across sectors such as aerospace, automotive, energy, climate modeling, and materials science. Panelists will share insights into how AI is accelerating simulations, enabling real-time optimization, and solving inverse problems that were previously intractable.

Lucas Nissenbaum – Instituto de Matemática Pura e Aplicada (IMPA) – (Google Scholar) – Panel Moderator
Lucas Nissenbaum is a Project Scientist at IMPA's Center for Projects and Innovation and the institution's Manager of Technological Projects. He holds a Master's and PhD in Electrical Engineering and Computer Science from MIT. His work focuses on developing academic–industrial collaborations in applied mathematics, with a particular emphasis on data science and machine learning. Currently, he leads a collaboration with Petrobras that aims to use physics-informed neural networks to solve partial differential equations, with a specific focus on seismic inversion applied to geophysical data. He is also developing a project for the automated reading of answer sheets from the Brazilian Public School Mathematics Olympiad (OBMEP), applying computer vision and image processing methods.

Jay Lee – University of Maryland – (Website, Google Scholar) – Panelist – to be confirmed
Dr. Jay Lee is Clark Distinguished Chair Professor and Director of Industrial AI Center in the Mechanical Engineering Dept. of the Univ. of Maryland College Park. His current research is focused on developing non-traditional machine learning including transfer learning, domain

adaptation, similarity-based machine learning, stream-of-x machine learning, as well as industrial large knowledge model (ILKM), etc. In addition, he is leading the Data Foundry which consists of over 100 diversified industrial datasets including semiconductor manufacturing, jet engines, wind turbine, EVs, high speed train, machine tools, robots, medical TBI, etc. for industrial AI talent development. These datasets are also used to rapidly develop and validate Industrial AI system with scalable and systematic approaches. In addition, he is developing an "AI Factory" initiative to establish an Industrial AI Open Lab. to develop new breed of industrial AI engineers.

John R. Smith – IBM Research – (Website, Google Scholar) – Panelist – CONFIRMED
Dr. John R. Smith is IBM Fellow and Head of AI for Math and Science at IBM T. J. Watson Research Center. Dr. Smith received Ph.D. in Electrical Engineering, Columbia University, 1997, where he was awarded the Eliahu I. Jury prize for outstanding thesis in signal processing. Dr. Smith has led research at IBM in diverse areas spanning computer vision, speech, language, multimedia, and scientific discovery in domains such as chemistry and materials, healthcare and life sciences, and climate and sustainability. Dr. Smith has published in many top journals and conferences (h-index = 88, i10-index = 350, #citations = 34K) and has been awarded more than 100 patents. Dr. Smith has received ACM SIGMM Award for Outstanding Technical Contributions to Multimedia Computing, Communications and Applications and is Fellow of IEEE.

Pablo Javier Blanco – Laboratório Nacional de Computação Científica (LNCC/MCTI) – (Website, Google Scholar) – Panelist – CONFIRMED
Dr. Pablo J. Blanco received his PhD in Computational Modeling from the LNCC (2008, Brazil). He is a Senior Permanent Researcher at the National Laboratory for Scientific Computing (LNCC/MCTI, Brazil) and Head of the HeMoLab (Hemodynamics Modeling Laboratory) group at the LNCC. Dr. Blanco was an Affiliated Member of the Brazilian Academy of Sciences (2014-2018), and is currently a member of the World Council of Biomechanics and co-PI of the National Institute of Science and Technology in Medicine Assisted by Scientific Computing (INCT-MACC). He has been awarded the Productivity Fellowship Level 1B (CNPq) and the Scientist of our State Fellowship (FAPERJ). Dr. Blanco has published over 130 refereed papers in international scientific journals (H-index = 32, with more than 3,100 citations, according to Scopus). His primary research focuses on modeling and simulating the cardiovascular system, computational hemodynamics, and multi-scale modeling techniques in multi-physics systems, with applications in both basic science and the medical field. Dr. Blanco is also responsible for leading the HeMoLab group through the construction of the ultimate cardiovascular simulator known as the ADAVN (Anatomically Detailed Arterial-Venous Network) model.

Pavel Dimitrov – NVIDIA – (Website) – Panelist – tentative
Pavel Dimitrov is a senior solutions architect at NVIDIA. He helps energy partners and customers discover what's possible with accelerated computing and, most recently, through the use of physics-ML. More broadly, he likes to learn math and to apply it to other fields of knowledge. Prior to NVIDIA, Pavel spent over 12 years in oil and gas collaborating with scientists and engineers on research and development projects resulting in 16 granted patents. He holds a Ph.D. in Computer Science from Yale University.

6 ORGANIZERS

Eduardo Soares – IBM Research – `eduardo.soares@ibm.com` – (Website, Google Scholar)
Eduardo Soares is a Senior Research Scientist at IBM Research, specializing in the design and pre-training of large multimodal foundation models for scientific applications. His research focuses on integrating data-driven learning to model complex dynamical behaviors across scientific domains. His broader interests include representation learning, scientific simulation, and large-scale AI architectures for advancing scientific discovery. Previously, he contributed to the development of explainable AI techniques for autonomous systems and trustworthy deep learning methodologies. Eduardo earned his Ph.D. in Computer Science from Lancaster University, where his research received the Best Doctoral Dissertation Award from the International Neural Networks Society in 2022.

Daniel Yukimura – Instituto de Matemática Pura e Aplicada (IMPA) – `yukimura@impa.br` – (Google Scholar)
Daniel Yukimura is a Project Scientist at the Center for Projects and Innovation at IMPA (Brazil), where he develops academic–industrial collaborations that apply advanced mathematics and AI to

real-world challenges. His research focuses on simulation methods and machine learning for scientific and engineering problems, with contributions spanning particle filtering techniques and the theoretical analysis of neural networks. He is currently working on physically informed machine learning for partial differential equations, with applications to seismic inversion in collaboration with Petrobras.

Nara Bobko – IMPA Tech – `nara.bobko@impatech.edu.br` – (Website, Google Scholar)
Nara Bobko is the Academic Manager at IMPA Tech, the undergraduate program of IMPA (Brazil), where she also serves as a professor and coordinator of the Bachelor's in Mathematics of Technology and Innovation. Her research interests focus on mathematical modeling applied to biological phenomena, particularly population dynamics and infectious disease modeling.

Arthur Bizzi – EPFL – `arthur.coutinhobizzi@epfl.ch` – (Website, Google Scholar)
Arthur Bizzi is a postdoctoral Research Associate at the École Polytechnique Fédérale de Lausanne (EPFL), at the Intelligent Maintenance and Operations Systems (IMOS) laboratory. His research interests lie in the application of numerical and analytical methods from the theory of dynamical systems to the development of Physics-Informed architectures for scientific machine learning.

Siddhartha Mishra – ETH Zürich – `smishra@math.ethz.ch` – (Website, Google Scholar)
Siddhartha Mishra is a chair Professor of Applied Mathematics at ETH Zurich, where he heads the Computational and Applied Mathematics Laboratory (CAMLab) and the Seminar for Applied Mathematics. He is also the Director of Computational Science Zurich, a core faculty member of the ETH AI Center and a steering committee member of the Swiss National AI institute. His research interests lie in the fields of numerical analysis, scientific computing and machine learning/AI with applications to different fields of science and engineering including fluid dynamics, astrophysics, climate science, geophysics and biology. He is an elected Fellow of the European Academy of Sciences and a recipient of the Collatz Prize of ICIAM, the Dahlquist Prize of SIAM, and the von Mises Prize of GAMM, among others. He was also an invited speaker at the International Congress of Mathematicians (ICM) in Rio in 2018.

Elisa Serioli – Dallara – `e.serioli@dallara.it` – (Website
Elisa Serioli is the Head of CFD Methodology at Dallara, where she leads the development and implementation of advanced Computational Fluid Dynamics techniques. Her work focuses on optimizing aerodynamic performance and ensuring the high-fidelity application of CFD tools in cutting-edge motorsport and engineering projects. She is a recognized leader in leveraging simulation to drive innovation and performance gains within the automotive industry.

Ana Muller – Petrobras – `anapmuller@petrobras.com.br` – (Google Scholar)
Ana Paula Muller has worked in the Geophysical Technology Department of Petrobras Exploration since 2014, developing algorithms for migration velocity analysis and FWI, fundamental tools that help ensure a high-quality seismic image and reduce uncertainty. Her research integrates deep learning methods into the velocity model-building flow, developing artificial intelligence technologies for Petrobras's seismic processing algorithms.

PROGRAM CHAIRS

The Program Chairs will oversee the review process, coordinate with the organizing team, and ensure fairness, diversity, and high scientific quality across all submissions. They will also liaise with the ICLR Workshop Chairs and maintain the schedule and logistics of accepted talks and posters.

Emilio Vital Brazil – IBM Research – `evital@br.ibm.com` (Website, Google Scholar)
Emilio Vital Brazil is Senior Research Scientist and Manager of the AI-Assisted Decision-Making Group at IBM Research Brazil. He holds a Ph.D. in Applied Mathematics from IMPA and works on AI, knowledge representation, and risk characterization for complex decision-making. Emilio has published over 40 papers in venues such as Communications Chemistry and npj Artificial Intelligence and is inventor on several U.S. patents in machine learning and visual analytics. His recent research includes foundation models for material discovery and multi-view architectures for molecular property prediction.

Renato Cerqueira – Behring Institute of AI at PUC-Rio – `rcerq@puc-rio.br` (Website, Google Scholar)
Renato Cerqueira is the Director of the Behring Institute for Artificial Intelligence at PUC-Rio, whose

mission is to align research and the advancement of AI technologies with topics of great societal impact. Before taking on this role, he was a Senior Research Manager at IBM Research Brazil, where he led the Knowledge-centric Systems group, investigating new technologies for Human-AI co-creation and scalable knowledge management in processes such as scientific discovery, decision-making under uncertainty, and data interpretation. In collaboration with various clients and external partners, Renato and his team explored the application of these research efforts in areas such as Materials Design, Geosciences, and Finance.

PROGRAM COMMITTEE

The Program Committee will be coordinated by the program chairs, with active support from the organizers, invited speakers, panelists, and institutional ambassadors. Together, they will recruit reviewers among researchers, postdoctoral fellows, and Ph.D. students from diverse institutions. We expect to gather approximately **100 reviewers across more than 20 institutions**, ensuring broad representation across regions, genders, and career stages. Each paper will receive **at least three reviews**, and the review load will not exceed three papers per reviewer. This distributed model guarantees constructive and detailed feedback while maintaining fairness and inclusivity in the evaluation process.

REFERENCES

Henrik Bruus and Karsten Flensberg. *Many-body quantum theory in condensed matter physics: an introduction*. Oxford university press, 2004.

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*, 37(12): 1727–1738, 2021.

Jingrun Chen, Xurong Chi, Zhouwang Yang, et al. Bridging traditional and machine learning-based algorithms for solving pdes: the random feature method. *J Mach Learn*, 1(3):268–298, 2022.

Chun-Wun Cheng, Jiahao Huang, Yi Zhang, Guang Yang, Carola-Bibiane Schönlieb, and Angelica I Aviles-Rivero. Mamba neural operator: Who wins? transformers vs. state-space models for pdes. *arXiv preprint arXiv:2410.02113*, 2024.

Lawrence C Evans. *Partial differential equations*, volume 19. American mathematical society, 2022.

A Fasoli, S Brunner, WA Cooper, JP Graves, P Ricci, O Sauter, and L Villard. Computational challenges in magnetic-confinement fusion physics. *Nature Physics*, 12(5):411–423, 2016.

Evangelos Galaris, Gianluca Fabiani, Ioannis Gallos, Ioannis Kevrekidis, and Constantinos Siettos. Numerical bifurcation analysis of pdes from lattice boltzmann model simulations: a parsimonious machine learning approach. *Journal of Scientific Computing*, 92(2):34, 2022.

YiMing Gao, Bing Wang, Jingyi Lu, and Zhou Tian. Ml-pinn: A memory-efficient physics-informed mamba-lstm network for fast and accurate pde solving. *Neurocomputing*, pp. 131446, 2025.

Maximilian Herde, Bogdan Raonic, Tobias Rohner, Roger Käppeli, Roberto Molinaro, Emmanuel de Bézenac, and Siddhartha Mishra. Poseidon: Efficient foundation models for pdes. *Advances in Neural Information Processing Systems*, 37:72525–72624, 2024.

George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.

Nikola Kovachki, Samuel Lanthaler, and Siddhartha Mishra. On universal approximation and error bounds for fourier neural operators. *Journal of Machine Learning Research*, 22(290):1–76, 2021.

Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to pdes. *Journal of Machine Learning Research*, 24(89):1–97, 2023.

Mark Lundstrom. Fundamentals of carrier transport, 2nd edn. *Measurement Science and Technology*, 13(2):230–230, 2002.

Giovanni Manfredi, Paul-Antoine Hervieux, and Jérôme Hurst. Fluid descriptions of quantum plasmas. *Reviews of Modern Plasma Physics*, 5(1):7, 2021.

Chuizheng Meng, Sam Griesemer, Defu Cao, Sungyong Seo, and Yan Liu. When physics meets machine learning: A survey of physics-informed machine learning. *Machine Learning for Computational Science and Engineering*, 1(1):20, 2025.

Tim N Palmer. *The Primacy of Doubt: From climate change to quantum physics, how the science of uncertainty can help predict and understand our chaotic world*. Oxford University Press, 2022.

Christopher J Pethick and Henrik Smith. *Bose–Einstein condensation in dilute gases*. Cambridge university press, 2008.

Pierre Sagaut and Claude Cambon. *Homogeneous turbulence dynamics*, volume 15. Springer, 2008.

Junhong Shen, Tanya Marwah, and Ameet Talwalkar. Ups: Efficiently building foundation models for pde solving via cross-modal adaptation. *arXiv preprint arXiv:2403.07187*, 2024.

Jingmin Sun, Yuxuan Liu, Zecheng Zhang, and Hayden Schaeffer. Towards a foundation model for partial differential equations: Multioperator learning and extrapolation. *Physical Review E*, 111(3): 035304, 2025.

Sifan Wang, Hanwen Wang, and Paris Perdikaris. Improved architectures and training algorithms for deep operator networks. *Journal of Scientific Computing*, 92(2):35, 2022.

Florian Wiesner, Matthias Wessling, and Stephen Baek. Towards a physics foundation model. *arXiv preprint arXiv:2509.13805*, 2025.