

---

# ICLR 2025 Workshop Proposal: The 2nd Workshop on Foundation Models for Science: Real-World Impact and Science-First Design

---

## 1. Workshop Description

**Title:** *The 2nd Workshop on Foundation Models for Science: Real-World Impact and Science-First Design*

**Website:** <https://fm-science.github.io/>

Scientific foundation models should be built **for science**, not for generic AI tastes or leaderboard prestige [8, 11]. This workshop centers *problem-driven design: models that measurably advance real scientific inquiries*, e.g., forecasting extreme climate events, accelerating materials discovery, understanding biological mechanisms, co-developed with domain experts and validated against field data, experiments, and downstream impact [4, 5, 7].

We argue that **foundation models for science must be built differently** from language and vision [2, 10]. Scientific data are physical, causal, spatiotemporal, and often scarce or biased; objectives must reflect mechanistic fidelity, not just predictive accuracy. This calls for *scientific priors and constraints, robust uncertainty quantification (UQ), and architectures that natively handle multi-modality (e.g., grids, meshes, spectra, time series, point clouds, text, images, code)*. It also demands tight integration with classical scientific tools (simulators, PDE solvers, optimization and inference engines, and HPC workflows) to yield hybrid systems that are faster, more accurate, and more trustworthy.

We will highlight opportunities and hard problems **unique to science**: enforcing conservation laws and symmetries; learning across vast spatial and temporal scales; representing extreme events and tipping points; calibrating and validating UQ; and developing evaluation protocols that reward mechanistic insight and actionable reliability [6, 3, 9, 12, 1]. The goal is a roadmap for building, training, and deploying scientific foundation models that accelerate discovery while respecting the structure of the natural world. We aim to bring together experts from foundation models and scientific problems, spur discussions, and foster collaborations on broad and transformative questions:

### 1. Real-World Impact

- *Problem-driven progress*: Demonstrate measurable gains on urgent scientific tasks (e.g., new scientific discovery, forecasting, prediction quality), not just leaderboards.
- *Reusability, adaptation, and trustworthy evaluation*: Pretrained models should transfer across instruments/labs/regimes with lightweight adaptation (fine-tuning, adapters, retrieval, prompting), deliver calibrated UQ grounded in scientific facts, and withstand rare-event stress tests.
- *Accelerating the scientific loop*: Close the loop of “hypothesis, design, experiment/simulation, analysis,” reducing cycle time and cost.
- *Operations across scales, extremes, and production*: Multi-scale (spatial/temporal) coupling with adaptive resolution and stable long-horizon rollouts; portable HPC/edge deployment with monitoring, diagnostics, and safeguards.

## 2. Science-First Design

- *Scaling laws for science*: Are data/model/compute scaling trends and training curricula fundamentally different from NLP/vision? What pretraining corpora best transfer?
- *Physical and scientific constraints*: Embed conservation laws, symmetries, invariants, and causal structure during training and inference (e.g., constrained decoding, differentiable physics).
- *Integration with classical tools*: Hybridize with simulators, PDE solvers, optimizers, and Bayesian inverse methods; quantify speed-accuracy trade-offs (emulators/surrogates).
- *Failure modes & diagnostics*: Systematically identify when/why models fail; interpretability for mechanism discovery; dataset/model audits and reproducibility.

**Scientific Domains.** We invite paper submissions from various scientific domains, including but not limited to Quantum Mechanics (e.g., nuclear fusion), Small Molecules, Biomedicine (e.g., proteins, biosequences, virtual screening), Materials Science (e.g., batteries, chemical synthesis), Earth Science, and Computational Science (e.g., PDEs, forecasting). Applications-driven submissions focusing on AI-for-Science and Scientific Machine Learning (SciML) are also highly encouraged.

**Tiny/Short Paper Track** In alignment with ICLR 2026 guidance, we will host a dedicated *tiny/short paper track* to encourage the submission of early-stage, high-potential ideas that may not yet be mature enough for full-length papers. Submissions to this track will undergo a light but fair review process focused on novelty, clarity, and potential impact. To ensure integrity and originality, **AI-generated papers will not be permitted for this track**. Selected tiny papers will be presented as posters or short talks to foster interactive discussion and community feedback.

## 2. Workshop Schedule and Logistics

Tentative important dates for paper submissions (anywhere on earth):

- Abstract Submission Deadline: January 28, 2026
- Paper Submission Deadline: January 30, 2026
- Review Bidding Period: January 31 - February 7, 2026
- Reviewer Deadline: February 28, 2026
- Acceptance/Rejection Notification Date: March 1, 2026
- Import Workshop Program and Accepted Papers to iclr.cc: March 11, 2026

Rio de Janeiro Time (GMT-3)	Event	Rio de Janeiro Time (GMT-3)	Event
8:55-9:00	Opening Remarks	13:30-14:10	Invited Talk 4
9:00-9:40	Invited Talk 1	14:15-14:55	Invited Talk 5
9:45-10:25	Invited Talk 2	15:00-15:30	Poster
10:25-10:55	Poster	15:35-16:15	Invited Talk 6
11:00-11:40	Invited Talk 3	16:20-17:00	Invited Talk 7
11:40-12:10	Contributed Talks	17:00-17:30	Contributed Talks
12:10-13:30	Lunch	17:30-17:35	Closing Remarks

### 3. Participation and Accessibility Plan.

At ICLR 2026, we anticipate receiving more than 100 paper submissions, with an acceptance rate of around 25%. Of the accepted papers, 16% will be invited for oral presentations.

**Attendees.** We expect 200-300 in-person attendees. Our goal is to attract a diverse range of participants, including: 1) entry-level graduate students, through the provision of educational materials; 2) early-career researchers, by offering a variety of research topics; and 3) established domain leaders, through opportunities for interdisciplinary communication and potential collaborations.

**Virtual Access.** We are committed to ensuring accessibility for those unable to attend physically. All accepted talks, posters, and panel sessions will have their slides and (where permitted) recordings posted on the workshop website. We will also maintain an active online presence (via the workshop site and social channels) to share materials, summaries, and key discussions with the broader community.

**Audience-Building and Outreach.** To attract a diverse and engaged audience, we will promote the workshop through multiple channels, including academic mailing lists (e.g., ML-News) and social media (Twitter/X, LinkedIn). We will also coordinate with ICLR and related workshops to cross-advertise.

**Special Requirements & Technical Needs.** No special needs other than the standard AV setup.

**Previous Attendance.** Our [1st workshop at NeurIPS 2024](#) received 100+ paper submissions and 100+ in-person attendees.

### 4. Invited Speakers (A-Z by Last Name)

This workshop has devoted enormous efforts to ensure the demographic diversity of speakers (3 female, 4 male). All speakers and organizers plan to attend the workshop in person. All speakers have **confirmed** to talk at our workshop.

[Steven Brunton](#): Professor, University of Washington.

[Aditi Krishnapriyan](#): Assistant Professor, UC Berkeley.

[Michael Mahoney](#): Group Lead, LBNL; Vice President, ICSI; Professor, UC Berkeley.

[Mahdi Soltanolkotabi](#): Professor, USC; Director, USC Center on AI Foundations for Science.

[Yuyang \(Bernie\) Wang](#): Principal Scientist, AWS AI.

[Rebecca Willett](#): Professor and Director of AI, Data Science Institute, University of Chicago.

[Rose Yu](#): Associate Professor, UC San Diego

### 5. Organizers & Program Committee

This workshop strives to ensure the demographic diversity of **confirmed** organizers (Appendix A):

[Wuyang Chen](#) (Simon Fraser Univ.), [Yongji Wang](#) (NYU), [Benjamin Erichson](#) (ICSI, and LBNL), [Laurence Perreault-Levasseur](#) (Univ. Montreal), [Bo Li](#) (UIUC), [Damian Borth](#) (University of St. Gallen), [Swarat Chaudhuri](#) (UT Austin).

We have also confirmed a list of Program Committee members (Appendix B), so that each paper receives 3 reviews, and no reviewer is committed to reviewing more than 3 papers.

## 6. Conflict of Interest (COI) Policy

To ensure fairness and transparency in the review and selection process, the organizers will **adhere to a strict conflict of interest (COI) policy**.

1. No organizer will handle or review submissions from their own institution, close collaborators, or any work with which they have a personal or professional conflict.
2. Organizers will not give talks or present papers during the workshop to avoid any perception of self-promotion.
3. The reviewing process will be conducted using double-blind principles, and assignments will be managed to prevent COI violations.
4. All program committee members and reviewers will be instructed to declare potential conflicts, and submissions with COIs will be reassigned to independent reviewers.

## 7. Relevant Workshops

This will be the 2nd version of our workshop. Our [1st workshop was at NeurIPS 2024](#). Two related directions in machine learning conferences.

- **Foundation Models.** Many recent workshops focused on discussing applications and properties of foundation models. This includes the “Foundation Models for Decision Making” workshop series ([Neurips 2023](#), [Neurips 2022](#)), “[Robustness of Zero/Few-Shot Learning in Foundation Models \(R0-FoMo\)](#)” ([Neurips 2023](#)), “[Distribution Shifts \(DistShift\) New Frontiers with Foundation Models](#)” ([Neurips 2023](#)), “[Federated Learning in the Age of Foundation Models](#)” ([Neurips 2023](#)), “[I Can’t Believe It’s Not Better \(ICBINB\): Failure Modes in the Age of Foundation Models](#)” ([Neurips 2023](#)), “[ES-FoMo: Efficient Systems for Foundation Models](#)” ([ICML 2023](#)), “[Mathematical and Empirical Understanding of Foundation Models \(ME-FoMo\)](#)” ([ICLR 2023](#)).
- **AI-for-Science.** Scientific problems become key focuses in the academia, stirring active discussion over methodologies and applications. Prominent examples are the “Machine Learning and the Physical Sciences” workshop series ([Neurips 2024](#), [Neurips 2023](#), [Neurips 2022](#), [Neurips 2021](#)), the “AI for Science” workshop series ([Neurips 2024](#), [Neurips 2023](#), [Neurips 2022](#), [ICML 2022](#), [Neurips 2021](#)), and the “[AI for Earth and Space Science](#)” ([ICLR 2022](#)).

In comparison, our workshop will emphasize how foundation models will change the game in the scientific domain, as the large scaling in both data, model, and training is proven to be key in qualitative changes in CV and NLP. Our workshop will provide timely discussions about potentials and opportunities in this research direction.

## A. Organizers

**Wuyang Chen** [Simon Fraser University] (wuyang@sfu.ca) [Google Scholar](#) | [Website](#) **Bio:** Dr. Wuyang Chen is an Assistant Professor in Computing Science at Simon Fraser University. Previously, Dr. Chen was a postdoc researcher in Statistics at the University of California, Berkeley, worked with Professor Michael Mahoney. He obtained his Ph.D. from the ECE Department at UT Austin in 2023, under the supervision of Professor Atlas Wang. Dr. Chen’s research focuses on the theoretical understanding of deep learning, with applications in foundation models, AutoML, computer vision, natural language processing, and addressing scientific problems. He published papers on CVPR, ECCV, ICLR, ICML, Neurips, etc. Dr. Chen’s work on training-free neural architecture design was highlighted as the "Featured Advances in Artificial Intelligence" in the National Science Foundation (NSF) newsletter in 2022. Dr. Chen co-organized the 1st workshop on Foundation Models for Science in NeurIPS 2024, and UG2+ workshop in CVPR.

**Yongji Wang** [New York University] (yw8211@nyu.edu) [Google Scholar](#) **Bio:** Yongji Wang is a postdoctoral associate at the Courant Institute of Mathematical Sciences at New York University, working with Prof. Tristan Buckmaster. He is also a visiting postdoctoral researcher in the Department of Geophysics at Stanford, working with Prof. Ching-Yao Lai. He received Ph.D. in Civil and Environmental Engineering from MIT and master’s degree in Applied Mathematics (Part III of the Mathematical Tripos) from the University of Cambridge with Distinction. His research focuses on using neural networks for multiscale inverse problems, from inferring ice-shelf rheology to finding self-similar blow-up solutions to fluid equations. He also focuses on developing techniques that enable physics-informed machine learning to reach machine precision. Dr. Wang co-organized the 1st workshop on Foundation Models for Science in NeurIPS 2024.

**N. Benjamin Erichson** [International Computer Science Institute] (erichson@icsi.berkeley.edu) [Google Scholar](#) | [Website](#) **Bio:** Ben Erichson is a senior research scientist at the International Computer Science Institute (ICSI), and a research scientist in Computing Sciences at Lawrence Berkeley National Laboratory. His research focuses on the intersection of deep learning and dynamical systems to develop more robust neural network architectures for scientific learning. Previously, he was a tenure-track Assistant Professor of Artificial Intelligence and Data-driven Modeling in the School of Engineering at the University of Pittsburgh. He completed postdoctoral research in the Department of Statistics at UC Berkeley, working with Michael Mahoney, and in the Department of Applied Mathematics at the University of Washington working with Nathan Kutz and Steven Brunton. Ben earned his PhD in Statistics from the University of St Andrews in Dec. 2017.

**Laurence Perreault-Levasseur** [University of Montreal] (laurence.perreaultlevasseur@gmail.com) [Google Scholar](#) | [Website](#) **Bio:** Prof. Laurence Perreault-Levasseur is the Canada Research Chair in Computational Cosmology and Artificial Intelligence. She is an Assistant Professor at the University of Montreal, an Associate Member of Mila, and a Visiting Scholar at the Center for Computational Astrophysics in NYC, where she conducts research at the intersection of machine learning and cosmological data analysis. She specializes in generative modeling and high-dimensional statistical inference problems in cosmology and astrophysics. Prof. Perreault-Levasseur regularly publishes in interdisciplinary venues such as Nature, NeurIPS and ICML, as well as in specialised astrophysical journals (ApJ, MNRAS, etc). Her previous experience in the organization of interdisciplinary events includes multiple workshops at main ML conferences, international conferences, and hackathons at the intersection of ML and astrophysics over the past 5 years, averaging 2 per year.

**Bo Li** [University of Illinois at Urbana-Champaign] (lbo@illinois.edu) [Google Scholar](#) | [Website](#) **Bio:** Dr. Bo Li is an Associate Professor in the Siebel School of Computing and Data Science at the University of Illinois at Urbana-Champaign. She is the recipient of the IJCAI Computers and Thought Award, Alfred P. Sloan Research Fellowship, NSF CAREER Award, AI’s 10 to Watch, MIT Technology Review TR-35 Award, Dean’s Award for Excellence in Research, C.W. Gear Outstanding Faculty Award, Intel Rising Star Award, Symantec Research Labs Fellowship, Rising Stars in EECS, Research Awards from Tech companies such as Amazon, Meta, Google, Intel, MSR, eBay, IBM, and best paper awards at top machine learning and security conferences. Her research focuses on theoretical and practical aspects of trustworthy machine learning, at the intersection of machine learning, security, privacy, and game theory. She has designed several scalable frameworks for robust learning and privacy-preserving data publishing systems. Her work has been featured by major publications and media outlets (Nature, Wired, Fortune, New York Times, etc.).

**Damian Borth** [University of St. Gallen] (damian.borth@unisg.ch) [Google Scholar](#) | [Website](#) **Bio:** Prof. Damian Borth is director of the Institute of Computer Science at the University of St. Gallen, where he holds a full professorship in Artificial Intelligence and Machine Learning (AIML). Previously, Damian was the founding director of the Deep Learning Competence Center at the German Research Center for Artificial Intelligence (DFKI) in Kaiserslautern, where he was also PI of the NVIDIA AI Lab at DFKI. Damian’s research focuses on representation learning of neural networks’ weight spaces.

His work has been awarded the ACM SIGMM Test of Time Award 2023, Google Research Scholar Award 2022, the NVIDIA AI Lab at NVIDIA GTC Europe 2016, the Best Paper Award at ACM ICMR 2012, the McKinsey Business Technology Award in 2011. Damian has served as Area Chair at NeurIPS, CVF/IEEE WACV and has organized several workshops in the past, such as the ICLR Workshop on Neural Network Weights as new Data Modality in 2025, the MMCommons workshops in 2015 and 2016, and the Deep Learning workshop at the International Supercomputing Conference (ISC) in 2017.

**Swarat Chaudhuri** [The University of Texas at Austin] (swarat@cs.utexas.edu) [Google Scholar](#) | [Website](#) **Bio:** Swarat Chaudhuri is a Professor of Computer Science and the director of the Trishul laboratory at UT Austin. His research lies at the interface of programming languages, formal methods, and machine learning. He aims to develop a new class of intelligent systems that are reliable, transparent, and secure by construction and can solve reasoning-intensive tasks beyond the scope of contemporary AI. Prof. Chaudhuri received a bachelor’s degree in computer science from the Indian Institute of Technology, Kharagpur (2001), and a Ph.D. in computer science from the University of Pennsylvania (2007). He has received the NSF CAREER award, the ACM SIGPLAN John Reynolds Dissertation award, the Morris and Dorothy Rubinoff Dissertation award from the University of Pennsylvania, Meta and Google Research awards, and several ACM SIGPLAN and SIGSOFT distinguished paper awards. He serves on the editorial boards of ACM Transactions on Programming Languages and Systems and Transactions on Machine Learning Research. He served as a Program Chair for CAV 2016 and ICLR 2024.

## B. Program Committee Members.

We confirm a list of Program Committee members who will serve as reviewers or area chairs for paper submissions to our workshop. We will ensure a sufficient number of reviewers so that each paper receives three reviews and no one is committed to reviewing more than three papers. We will continue to secure more reviewers before the paper submission deadline.

**Program Committee (97 members):** Siavash Ameli (UC Berkeley), Shashank Subramanian (LBNL), Ryan Theisen (Harmonic Discovery), Arnur Nigmatov (LBNL), Peter Harrington (LBNL), Jialin Song (SFU), Yuqiu Liu (SFU), Mauricio Soroco (SFU), Oleg Balabanov (UC Berkeley), Dejia Xu (UT Austin), Jonathan Liu (UT Austin), Junbo Li (UT Austin), Neel Bhatt (UT Austin), Greg Holste (UT Austin), Ajay Jaiswal (UT Austin), Atlas Wang (UT Austin), Wes Robbins (UT Austin), Scott Hoang (UT Austin), Ruisi Cai (UT Austin), Wenyan Cong (UT Austin), Runjin Chen (UT Austin), Yan Zheng (UT Austin), Ziwei Yang, Amin Totounferoush (Univ. Stuttgart), Konstantin Schürholt (Univ. St.Gallen), Krti Tallam (ICSI), Danielle Maddix Robinson (Amazon), Geoffrey Negiar, Ryan Theisen, Rui Wang (MIT), Luning Sun (LLNL), Han Gao (Harvard), Qingqing Zhao (Stanford), Wenbin Xu (LBNL), Xiaoyu Xie (Northwestern University), Rasmus Malik Hoeegh Lindrup (UC Berkeley), Konstantin Rusch (MIT), Lu Lu (Yale), Sifan Wang (Univ. Pennsylvania), Zhichao Wang (UCSD), Hengrui Luo (Rice), Max Daniels (MIT), Senwei Liang (LBNL), Fangzheng Sun (Amazon), Zhong Yi Wan (Google), Jiequn Han (Flatiron Institute), Yihang Gao (HKU), Zhengfa Bi (LBNL), Yingheng Tang (LBNL), Yaoqing Yang (Dartmouth), Yefan Zhou (Dartmouth), Jerry Liu (Stanford), Annan Yu (Cornell), Dongwei Lyu (UC Berkeley), Kareem Hegazy (ICSI), Rie Nakata (LBNL), Alejandro Queiruga (Google), John Cava (ASU), Wahid Bhimji (LBNL), Nori Nakata (LBNL), Cheng-Nan Liu (U. Utah), Garry Gao (ICSI), Yixiao Kang (Meta), Vinicius Mikuni (LBNL), Lanasse Francois (Simons Foundation), Mishra-Sharma Siddharth (IAIFI), Adam Alexandre (Mila/Univ. Montreal), Sotoudeh Hadi (Cambridge), Parker Liam (Flatiron Institute), Pettee Mariel (LBNL), Stone Connor (Mila/Univ. Montreal), Leung Henri (Univ. Toronto), Speagle Joshua (Univ. Toronto), Walmsley Mike (Univ. Toronto), Bowles Micah (Oxford), Ting Yuan-Sen (Australian National Univ.), Huertas-Company Marc (Instituto de Astrofísica de Canarias), Domínguez-Sánchez Helena (Centro de Estudios de Física del Cosmos de Aragón), Dia Noe (Mila/Univ. Montreal), Barth Michael (Mila/Univ. Montreal), Scaife Anna M. M. (Univ. Manchester), Pasquato Mario (Roma), Rhea Carter (Univ. Montreal), Mezcua Mar (Institute of Space Sciences), Charlie Cowen-Breen (MIT), Zongyi Li (Caltech), Nikola Kovachki (Nvidia), Facu Sapienza (Berkeley), Hannah Lu (MIT), Gong Chen (Dartmouth) Bryan Riel (Zhejiang Univ.), Chris Pedersen (NYU), Adam Subel (NYU), Johannes Müller (Max Planck), Marius Zeinhofer (Oslo Univ.), Konstantin Schürholt (St. Gallen).

## References

- [1] Bodnar, C., Bruinsma, W. P., Lucic, A., Stanley, M., Brandstetter, J., Garvan, P., Riechert, M., Weyn, J., Dong, H., Vaughan, A., et al. Aurora: A foundation model of the atmosphere. *arXiv preprint arXiv:2405.13063*, 2024.
- [2] Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- [3] Jakubik, J., Roy, S., Phillips, C., Fraccaro, P., Godwin, D., Zadrozny, B., Szwarcman, D., Gomes, C., Nyirjesy, G., Edwards, B., et al. Foundation models for generalist geospatial artificial intelligence. *arXiv preprint arXiv:2310.18660*, 2023.
- [4] Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.
- [5] Merchant, A., Batzner, S., Schoenholz, S. S., Aykol, M., Cheon, G., and Cubuk, E. D. Scaling deep learning for materials discovery. *Nature*, 624(7990):80–85, 2023.
- [6] Nguyen, T., Brandstetter, J., Kapoor, A., Gupta, J. K., and Grover, A. Climax: A foundation model for weather and climate. *arXiv preprint arXiv:2301.10343*, 2023.
- [7] Raissi, M., Perdikaris, P., and Karniadakis, G. E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.
- [8] Song, S. L., Kruft, B., Zhang, M., Li, C., Chen, S., Zhang, C., Tanaka, M., Wu, X., Rasley, J., Awan, A. A., et al. Deepspeed4science initiative: Enabling large-scale scientific discovery through sophisticated ai system technologies. *arXiv preprint arXiv:2310.04610*, 2023.
- [9] Stevens, S., Wu, J., Thompson, M. J., Campolongo, E. G., Song, C. H., Carlyn, D. E., Dong, L., Dahdul, W. M., Stewart, C., Berger-Wolf, T., et al. Bioclip: A vision foundation model for the tree of life. *arXiv preprint arXiv:2311.18803*, 2023.
- [10] Yuan, L., Chen, D., Chen, Y.-L., Codella, N., Dai, X., Gao, J., Hu, H., Huang, X., Li, B., Li, C., et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021.
- [11] Zhang, X., Wang, L., Helwig, J., Luo, Y., Fu, C., Xie, Y., Liu, M., Lin, Y., Xu, Z., Yan, K., et al. Artificial intelligence for science in quantum, atomistic, and continuum systems. *arXiv preprint arXiv:2307.08423*, 2023.
- [12] Zhu, X. X., Xiong, Z., Wang, Y., Stewart, A. J., Heidler, K., Wang, Y., Yuan, Z., Dujardin, T., Xu, Q., and Shi, Y. On the foundations of earth and climate foundation models. *arXiv preprint arXiv:2405.04285*, 2024.