Published papers in FY24

Machine Learning for Environmental and Earth Sciences

Fan, M., Liu, S., Lu, D., Gangrade, S., & Kao, S. (2024). Explainable machine learning model for multi-step forecasting of reservoir inflow with uncertainty quantification. Environmental Modelling & Software. https://doi.org/10.1016/j.envsoft.2023.105849

Fan, M., Liu, S., & Lu, D. (2023). Advancing subseasonal reservoir inflow forecasts using an explainable machine learning method. Journal of Hydrology: Regional Studies. https://doi.org/10.1016/j.ejrh.2023.101584

Fan, M., Zhang, L., Liu, S., Yang, T., & Lu, D. (2023). Investigation of hydrometeorological influences on reservoir releases using explainable machine learning methods. Frontiers in Water. https://doi.org/10.3389/frwa.2023.1112970

Gao, B., Coon, E. T., Thornton, P. E., & Lu, D. (2024). Improving the estimation of atmospheric vapor pressure using interpretable LSTM. Agricultural and Forest Meteorology. https://doi.org/10.1016/j.agrformet.2024.109907

Fan, M., Wang, H., Zhang, J., Hosseini, S., & Lu, D. (2024). Advancing spatiotemporal forecasts of CO2 plume migration using deep learning networks with transfer learning and interpretation analysis. International Journal of Greenhouse Gas Control. https://doi.org/10.1016/j.ijggc.2024.104061

Tayal, K., Renganathan A., and Lu, D., Improving Streamflow Predictions across CONUS by Integrating Advanced Machine Learning Models and Diverse Data. Environmental Research Letters. Doi: 10.1088/1748-9326/ad6fb7

Jamil, A., Rucker, D.. Lu, D., Cao, H., Brooks, S., and Carroll K., Comparison of machine learning and electrical resistivity arrays to inverse modeling for locating and characterizing subsurface targets, Journal of Applied Geophysics, 2024. https://doi.org/10.1016/j.jappgeo.2024.105493.

Wang, X., Liu, S., Tsaris, A., Choi, J.-Y., Aji, A., Fan, M., Zhang, W., Yin, J., Ashfaq, M., Lu, D., & Balaprakash, P. (2024). Oak Ridge Base Foundation Model for Earth System Predictability. International Journal of High Performance Computing Applications. (Accepted). https://arxiv.org/html/2404.14712v1

Deep Learning and Neural Networks

Tayal, K., Renganathan, A., & Lu, D. (2024). Improving streamflow predictions with vision transformers. ICLR Tackling Climate Change with Machine Learning Workshop. https://www.climatechange.ai/papers/iclr2024/4

Tayal, K., Renganathan, A., & Lu, D. (2024). FutureTST: When transformers meet future exogenous drivers. ICML Next Generation of Sequence Modeling Architectures Workshop. https://sites.google.com/view/ngsmworkshop/accepted-papers

Bororwiec, K., Lu, D., Chandan, V., Chatterjee, S., Ramuhalli, P., Tipireddy, R., Halappanavar, M., and Liu, F., Accelerating Scientific Simulations with Bi-Fidelity Weighted Transfer Learning, 23nd IEEE International Conference on Machine Learning and Applications. Doi: 10.1109/ICMLA58977.2023.00147, https://ieeexplore.ieee.org/document/10459994.

Yoo, P., Bhowmik, D., Mehta, K., Zhang, P., Liu, F., Lupo Pasini, M., & Irle, S. (2023). Deep learning workflow for the inverse design of molecules with specific optoelectronic properties. Scientific Reports, 13, 20031. https://doi.org/10.1038/s41598-023-45385-9

Zhang, E., Lyngaas, I., Chen, P., Wang, X., Igarashi, J., Huo, Y., Wahib, M., & Munetomo, M. (2024). Adaptive Patching for High-resolution Image Segmentation with Transformers. arXiv preprint arXiv:2404.09707. (Accepted to SC'24). https://arxiv.org/abs/2404.09707

Lyngaas, I., Meena, M. G., Calabrese, E., Wahib, M., Chen, P., Igarashi, J., Huo, Y., & Wang, X. (2024). Efficient Distributed Sequence Parallelism for Transformer-Based Image Segmentation. Electronic Imaging, 36, 1-7.https://library.imaging.org/admin/apis/public/api/ist/website/downloadArticle/ei/36/12/HPCI-199

Gurecky, W. (2024). Latent neural controlled differential equations for time series forecasting. Presented at the International Conference for Physics of Reactors (PHYSOR24), April 21st, 2024. Invited workshop presentation. URL with published slides: https://www.ans.org/meetings/physor2024/session/view-2266/ (accessible to ANS members).

Y Xue, J Zha, M Wahib, T Ouyang, X Wang, “Neural architecture search via similarity adaptive guidance”, in Applied Soft Computing, Page 111821, Elsevier. https://www.sciencedirect.com/science/article/abs/pii/S1568494624005957

Y Xue, W Tong, F Neri, P Chen, T Luo, L Zhen, X Wang, “Evolutionary architecture search for generative adversarial networks based on weight sharing” in *IEEE Transactions on Evolutionary Computation*, vol. 28, no. 3, pp. 653-667, June 2024. https://ieeexplore.ieee.org/document/10336909

AI for Materials Science and Chemistry

Lupo Pasini, M., Karabin, M., & Eisenbach, M. (2024). Transferring predictions of formation energy across lattices of increasing size. Machine Learning: Science and Technology, 5(2), 025015. https://doi.org/10.1088/2632-2153/ad3d2c

Lupo Pasini, M. (2024). First-principles data for solid solution niobium-tantalum-vanadium alloys with body-centered-cubic structures. Nature Scientific Data. (Accepted, preprint available at ChemRxiv) https://chemrxiv.org/engage/chemrxiv/article-details/66aa8490c9c6a5c07a9f6106

Fox, Z., & Ghosh, A. (2024). Active Causal Learning for Decoding Chemical Complexities with Targeted Interventions. Machine Learning: Science and Technology, 5 035056. https://iopscience.iop.org/article/10.1088/2632-2153/ad6feb

Ghosh, S., Tom, A., Dasgupta, D., Ghosh, A. and Wirth, B. (2024). Insights into Prismatic Loop Formation in Irradiated Fe-Cr Alloys from Hypothesis-Driven Active Learning and Causal Analysis. ACS Applied Energy Materials, 7(15) , 6123-6134. https://pubs.acs.org/doi/10.1021/acsaem.4c00485

Ghosh, A. (2024). Towards physics-informed explainable machine learning and causal models for materials research. Computational Materials Science, 233, 112740. https://doi.org/10.1016/j.commatsci.2023.112740

Ghosh, A., Gayathri, P., Shaikh, M., & Ghosh, S. (2024). Structural mode coupling in perovskite oxides using hypothesis-driven active learning. Journal of Physics: Materials, 7, 025014. https://doi.org/10.1088/2515-7639/ad3fea

Palanichamy, G., Ghosh, S., & Ghosh, A. (2023). Predictive design of hybrid improper ferroelectric double perovskite oxides. Chemistry of Materials, 36, 682. https://doi.org/10.1021/acs.chemmater.3c02067

Swamynadhan, M. J., Ghosh, A., & Ghosh, S. (2023). Design of high polarization low switching barrier hybrid improper ferroelectric perovskite oxide superlattices. Materials Horizons, 10(12), 5942-5949. https://pubs.rsc.org/en/content/articlelanding/2023/mh/d3mh01285a

Ghosh, A., Trujillo, D. P., Hazarika, S., Schiesser, E., Swamynathan, M. J., Ghosh, S., Zhu, J.-X., & Nakhmanson, S. (2023). Identification of novel organic polar materials: A machine learning study with importance sampling. APL Machine Learning, 1(4).https://pubs.aip.org/aip/aml/article/1/4/046115/2928653

Han, B., Savici, A. T., Cheng, Y. (2024). INSPIRED: Inelastic neutron scattering prediction for instantaneous results and experimental design. Computer Physics Communications, 304, 109288. https://doi.org/10.1016/j.cpc.2024.109288

Okabe, R., Chotrattanapituk, A., Boonkird, A., Andrejevic, N., Fu, X., Jaakkola, T. S., Song, Q., Nguyen, T., Drucker, N., Mu, S., Wang, Y., Liao, B., Cheng, Y., Li, M. (2024). Virtual node graph neural network for full phonon prediction. Nature Computational Science, 4, 522–531. https://doi.org/10.1038/s43588-024-00661-0

Linker, T. M., Krishnamoorthy, A., Daemen, L. L., Ramirez-Cuesta, A. J., Nomura, K., Nakano, A., Cheng, Y., Hicks, W. R., Kolesnikov, A. I., Vashishta, P. D. (2024). Neutron scattering and neural-network quantum molecular dynamics investigation of the vibrations of ammonia along the solid-to-liquid transition. Nature Communications, 15, 3911. https://doi.org/10.1038/s41467-024-48246-9

AI and High-Performance Computing

Choi, J. Y., Lupo Pasini, M., Zhang, P., Mehta, K., Liu, F., Bae, J., & Ibrahim, K. (2023). DDStore: Distributed Data Store for Scalable Training of Graph Neural Networks on Large Atomistic Modeling Datasets. In SC-W '23: Proceedings of the SC '23 Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis (pp. 941-950). https://doi.org/10.1145/3624062.3624171

Lupo Pasini, M. (2024). AI for Materials Design and Discovery Using Atomistic Scale Information [Industrial and Governmental Activities]. IEEE Computational Intelligence Magazine, 19(2), 13-14. https://doi.org/10.1109/MCI.2024.3365234

Lupo Pasini, M. (2024). A Perspective on Scalable AI on High-Performance Computing and Leadership Class Supercomputing Facilities [Industrial and Governmental Activities]. IEEE Computational Intelligence Magazine, 19(3), 6-8. https://doi.org/10.1109/MCI.2024.3402770

Mehta, K., Lupo Pasini, M., Irle, S., Yoo, P., Suter, F., Ganyushin, D., & Klasky, S. (2024). Scaling Ensembles of Data-Intensive Quantum Chemical Calculations for Millions of Molecules. In The 25th IEEE International Workshop on Parallel and Distributed Scientific and Engineering Computing (PDSEC 2024). San Francisco, California, USA. https://doi.org/10.1109/IPDPSW63119.2024.00175

Dash, S., Lyngaas, I. R., Yin, J., Wang, X., Egele, R., Ellis, J. A., Maiterth, M., Cong, G., Wang, F., & Balaprakash, P. (2024). Optimizing distributed training on frontier for large language models. In ISC High Performance 2024 Research Paper Proceedings (39th International Conference) (pp. 1-11). Prometeus GmbH. https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=10528939

Yin, J., Bose, A., Cong, G., Lyngaas, I., & Anthony, Q. (2024). Comparative Study of Large Language Model Architectures on Frontier. In 2024 IEEE International Parallel and Distributed Processing Symposium (IPDPS) (pp. 556-569). San Francisco, CA, USA. https://doi.org/10.1109/IPDPS57955.2024.00056

E Zhang, I Lyngaas, P Chen, X Wang, J Igarashi, Y Huo, M Wahib, M Munetomo, “Adaptive Patching for High-resolution Image Segmentation with Transformers”, The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC’24). https://arxiv.org/abs/2404.09707

AI for Planning and Optimization

Tayal, K., Renganathan, A., & Lu, D. (2024). Planning for floods and droughts: Intro to AI-driven hydrological modeling. ICLR Tackling Climate Change with Machine Learning Workshop. https://www.climatechange.ai/papers/iclr2024/77

Li, W., Liu, Y., Zhang, C., Womble, D., & Wang, C. (2024). Distributed Reinforcement Learning with GNN Integration for Network-Level Traffic Signal Control. ASCE International Conference on Transportation & Development (ASCE ICTD 2024), June 15-18, 2024. Atlanta, GA.

AI Theory

Chowdhury, A., & Ramuhalli, P. (2024). A Provably Accurate Randomized Sampling Algorithm for Logistic Regression. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 38, No. 10, pp. 11597-11605). https://arxiv.org/abs/2402.16326.

AI for Power Grid

Liu B., Zhang C., Dong J., Balaprakash P., Liu Y., and Eiffert B., “Enhancing Power Distribution System Resilience with Fusion-GNN: A Dynamic Graph Representation Learning Approach”, in IEEE Industrial Electronics Society (IECON), Chicago, IL, Nov. 3-6, 2024. (Accepted)

AI for health

Y Wang, L Zhen, J Zhang, M Li, L Zhang, Z Wang, Y Feng, Y Xue, X Wang, Z Chen, T Luo, R Siow Mong Goh, Y Liu, “MedNAS: Multi-Scale Training-Free Neural Architecture Search for Medical Image Analysis”, IEEE Transactions on Evolutionary Computation, Volume 28, Pages 668-681. https://ieeexplore.ieee.org/document/10391077

Published softwares in FY24

EDLSTM-UQ-Explainability

An explainable machine learning model for improved multi-step reservoir inflow forecasting with uncertainty quantification. Features:

- Encoder-decoder long short-term memory (ED-LSTM) network for multi-step forecasting
- SHapley Additive exPlanation (SHAP) for interpreting hydrometeorological factors' influence
- Novel uncertainty quantification method for assessing prediction trustworthiness

GitHub: https://github.com/patrickfan/EDLSTM-UQ-Explainability

Exponential Mechanism with Normalizing Flows

An implementation of the Exponential Mechanism (ExpM) for private optimization and machine learning using Normalizing Flows. Key aspects:

- Operationalizes ExpM for continuous sample spaces
- Uses auxiliary Normalizing Flow for approximate sampling from ExpM density
- Provides an elegant solution for private ML, bypassing inefficiencies of DPSGD

GitHub: https://github.com/bridgesra/expm_nf_mimic3_results_code

Scalable Real-Time Data Assimilation Framework

A scalable implementation of a real-time data assimilation (DA) framework for estimating turbulent atmospheric dynamics. Components:

- Ensemble score filter (EnSF)
- Vision Transformer (ViT) surrogate
- Coupled experiments for prediction
- Demonstrates strong and weak scaling on the Frontier supercomputer
- Supports large-scale ViT training up to 2.5B parameters

GitHub: https://github.com/jqyin/sqg_vit

AIRES-NODE

The software package provides simplified interfaces to apply Neural Ordinary Differential Equations (NODEs) and Neural Controlled Differential Equations (NCDEs) to multivariate time series. The models are robust to data with irregular sample intervals. In contrast to traditional LSTMs or RNNs, NODEs and NCDEs are continuous time models. Additionally, this package provides interfaces to build and fit Latent VAE-NCDEs, a new technique that leverages past work on latent neural ordinary differential equations to enable robust time series forecasting given noisy measurements. The VAE-NCDE is a generative time series model capable of probabilistic multivariate time series forecasting.

GitHub: https://code.ornl.gov/AIRES/aires-node
DOE CODE: https://doi.org/10.11578/dc.20240124.2

MINNERVVA linter

https://github.com/minnervva/torchdetscan

This is a tool for finding non-deterministic functions in PyTorch code. It is coupled with a testing tool (still in finalization steps) that tests PyTorch functions using a range of arguments to determine the sensitivity of any non-deterministic functions to specific inputs, and to verify that the PyTorch reproducibility documentation is correct, for any combination of software stack versions, and on various hardware resources. The testing tool was first used in its alpha version in the following paper, which was recently submitted to SC24 Correctness Workshop: https://arxiv.org/abs/2408.05148 . The repository containing all codes and data required to reproduce the SC24 paper can be found here:

https://github.com/minnervva/correctness , including an extensive readme with additional details and supplementary data to support the SC24 paper.

We have found that the documentation for PyTorch's reproducibility support is not always completely accurate, thus the need for the testing tool for full numerical assurance.

Together, these two tools make up our new multimodal reproducibility testing suite for PyTorch. The paper introducing this package, and the full package with documentation, is currently in preparation.

DDStore

Enabling efficient distributed data loading for distributed data parallelism. DDStore provides an in-memory distributed data store designed for large-scale GNN training on US-DOE supercomputing facilities. DDStore provides a hierarchical, distributed, data caching technique that combines data chunking, replication, low-latency random access, and high throughput communication. DDStore achieves near-linear scaling for training a GNN model using up to thousands of GPUs on the Frontier, Summit and Perlmutter supercomputers, and reaches up to a 6.15x reduction in GNN training time compared to state-of-the-art methodologies.

GitHub: https://github.com/ORNL/DDStore
DOE CODE: https://www.osti.gov/doecode/biblio/121881

INSPIRED

A graphic user interface for rapid prediction and calculation of inelastic neutron scattering (INS) spectra. Key modules:
1. Direct prediction of phonon density of states, 1D and 2D INS spectra from a symmetry-aware neural network
2. INS simulation based on existing DFT-calculated forces
3. INS simulation using pre-trained universal machine learning force fields

GitHub: https://github.com/neutrons/inspired
Zenodo: https://zenodo.org/records/11478889

Software Demos and Tutorials

- David M. Rogers. Collecting Atomic Datasets. ORNL Developer Blog, August 29, 2024. (https://developer.ornl.gov/collecting-atomic-datasets/)

- Max Lupo Pasini. Training Graph Neural Network Models at Scale on DOE Leadership Class Supercomputing Facilities. 2024 AI Summer School. July 25, 2024.
- Max Lupo Pasini, Jong Youl Choi, Pei Zhang and Kshitij Mehta. Material Property Prediction with Large-Scale GNNs. ORNL AI for Science Tutorial Series, March 28, 2024.
- Pei Zhang, Jong Youl Choi, David M. Rogers, and Max Pasini. Predicting Chemical and Material Molecular Properties With HydraGNN. ORNL Developer Blog, February 8, 2024. (https://developer.ornl.gov/predicting-chemical-and-material-molecular-properties-with-hydragnn)
- Max Lupo Pasini, Jong Youl Choi, and Pei Zhang. Scalable GNN Training using HPC & Supercomputing. Learning on Graphs Conference 2023. November 28th, 2023. (http://log2023.logconference.org/, Learning on Graphs Conference 2023 - Day 2 - Oral Presentations and Tutorial)
- William Gurecky. VAE-NCDE: a generative time series model for probabilistic multivariate time series forecasting. 2024 AI Summer School (https://ornl.github.io/events/ai-summer-school-2024/). July 26, 2024. URL to tutorial: https://code.ornl.gov/AIRES/aires-node-tutorials
- Gurecky, W. Latent neural controlled differential equations for time series forecasting. Presented at the International Conference for Physics of Reactors (PHYSOR24), April 21st, 2024. Invited workshop presentation. URL with published slides: https://www.ans.org/meetings/physor2024/session/view-2266/ (accessible to ANS members).
- Yongqiang Cheng, Inelastic Neutron Scattering Analysis Software: INSPIRED, SHIVER and Sunny (only the INSPIRED part of the tutorial), June 23, 2024. URL to tutorial: https://ceramics.org/event-subpage/acns-2024-tutorials/