Delft University of Technology
Lawrence Livermore National Laboratory

Project: Learning Preconditioners Using Neural Networks

Supervisor(s):
Alexander Heinlein (Numerical Analysis, EEMCS)
Rui Peng Li (Lawrence Livermore National Laboratory)

October 9, 2025

Topic

Artificial intelligence (AI) now augments high-performance computing (HPC) at U.S. Department of Energy (DOE) laboratories such as Lawrence Livermore National Laboratory (LLNL). Flagship Exascale systems—including Frontier, Aurora, and LLNL's El Capitan—pair deep learning with large-scale solvers to accelerate simulations and data analysis across materials, climate, and physics applications. These AI-enabled workflows combine data-driven surrogates with multiphysics solvers to support rapid design loops, online calibration, and adaptive experimental control.

Iterative solvers and preconditioning remain central topics in numerical analysis and scientific computing, yet multiscale heterogeneities, multi-physics coupling, and strong nonlinearities still strain classical techniques. This project, developed in close collaboration with LLNL researchers, investigates whether machine learning (ML) can provide guidance where analytical tools are limited.

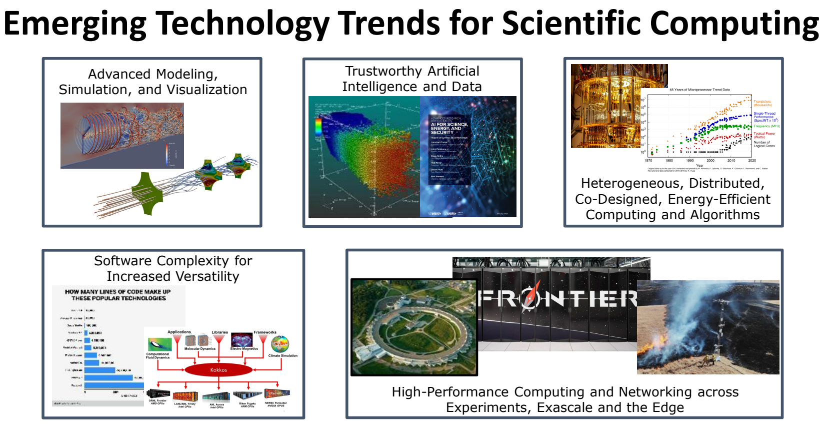

(a) AI initiatives across DOE laboratories.

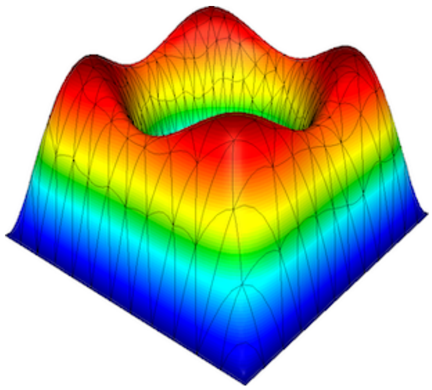

(b) Heterogeneous diffusion benchmark.

Figure 1: AI–HPC initiatives (a) and the heterogeneous diffusion benchmark (b) motivating learned preconditioners.

Preconditioning accelerates iterative solvers for large-scale systems, yet incomplete LU (ILU) factorization [4] and sparse approximate inverse (SPAI) methods [2, 3] depend on fragile heuristics to manage fill-in across heterogeneous regimes. We train message-passing graph neural networks (GNNs) to learn inverse mappings from matrix structure and governing physics so that data-driven preconditioners meet or surpass classical performance on partial differential equation (PDE) matrices while keeping physics-aware features and convergence diagnostics interpretable, building on recent operator-learning advances in graph networks [1, 5]. In the preconditioned system

$$M^{-1}Ax = M^{-1}b, \tag{1}$$

A denotes the typically sparse stiffness matrix, x the unknown state vector, and b the discretized forcing, while the neural network learns an operator M^{-1} that stays inexpensive to apply yet accelerates convergence. Because a sparse PDE matrix naturally defines a weighted graph over its degrees of freedom, GNNs offer a structured way to encode locality and multiscale coupling. Figure 1b shows the heterogeneous diffusion benchmark used to compare learned and traditional preconditioners.

Tasks

- Explore neural network-based approaches for learning preconditioners that exploit matrix structure, for example using graph neural networks (GNNs).
- Design loss functions that drive rapid convergence of iterative solvers for the preconditioned system (1).
- Implement and evaluate the approach in a finite element framework on representative PDE-based test problems.
- Optional: conduct HPC experiments on DelftBlue or LLNL systems.

Requirements

- Background in numerical linear algebra, iterative solvers, and partial differential equations (PDEs).
- Experience with machine learning frameworks (e.g., PyTorch, JAX) in Python or Julia.
- Optional: Familiarity with high-performance computing (HPC) workflows and automation for large-scale experiments.

Contact

Are you interested or do you have any questions? Send an email to Alexander Heinlein (a.heinlein@tudelft.nl) and/or Rui Peng Li (li50@llnl.gov).

References

[1] P. Battaglia, J. B. C. Hamrick, V. Bapst, A. Sanchez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner, C. Gulcehre, F. Song, A. Ballard, J. Gilmer, G. E. Dahl, A. Vaswani, K. Allen, C. Nash, V. J. Langston, C. Dyer, N. Heess, D. Wierstra, P. Kohli, M. Botvinick, O. Vinyals, Y. Li, and R. Pascanu. Relational inductive biases, deep learning, and graph networks. *arXiv*, 2018. URL https://arxiv.org/pdf/1806.01261.pdf.

[2] E. Chow. A priori sparsity patterns for parallel sparse approximate inverse preconditioners. *SIAM Journal on Scientific Computing*, 21(5):1804–1822, 2000. doi: 10.1137/S106482759833913X. URL https://doi.org/10.1137/S106482759833913X.

[3] M. J. Grote and T. Huckle. Parallel preconditioning with sparse approximate inverses. *SIAM Journal on Scientific Computing*, 18(3):838–853, 1997. doi: 10.1137/S1064827594261242.

[4] Y. Saad. *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2 edition, 2003.

[5] S. Yusuf, J. E. Hicken, and S. Pan. *Constructing ILU Preconditioners for Advection-Dominated Problems Using Graph Neural Networks*. doi: 10.2514/6.2024-3613. URL https://arc.aiaa.org/doi/abs/10.2514/6.2024-3613.