M2 Stage: Parallel training algorithms for Graph Neural Networks (GNNs)

Research project

In weather and climate modeling, many types of unstructured or spatially irregular data can be represented using graphs, which naturally arise from sensor networks, irregular observation grids, or heterogeneous spatial interactions. Graph neural networks (GNNs) [4, 3] are well suited for processing such data. A GNN treats each spatial location or sensor as a node and uses edges to encode physical proximity or atmospheric dependencies. By iteratively passing and aggregating information across neighboring nodes, GNNs can capture complex spatial relationships and multiscale patterns, making them valuable for meteorological tasks such as forecasting, sensor fusion, and data assimilation.

Despite their effectiveness, GNNs often suffer from tedious and time-consuming training phases, particularly when dealing with large and complex graphs. To overcome this difficulty, this project aims to accelerate GNN training by leveraging domain decomposition methods (DDMs), which have been successfully used in scientific computing for performing parallel large-scale simulations of complex physical phenomena. In particular, we aim to extend the capabilities of traditional DDMs beyond solving partial differential equations (PDEs) by adapting them to the non-convex optimization problems underlying GNN training. The key idea is to exploit the graph's intrinsic connectivity. Namely, since nodes are not uniformly linked and connection strengths vary with graph weights, we aim to identify a suitable partitioning that groups strongly connected nodes within the same subdomain while keeping inter-partition links comparatively weak, thereby enabling an efficient and scalable decomposition for parallel GNN training.

During the internship, the student will contribute to the design of the domain-decomposition-based training algorithms [2, 1], implement them within existing machine-learning libraries, and evaluate their performance using GNNs with realistic datasets relevant to meteorological applications. Beyond its theoretical interest, this project has strong practical relevance: improving the training and scalability of GNNs could enable faster and more efficient analysis of large-scale weather predictions. The internship, therefore, offers a unique opportunity to contribute to cutting-edge research at the intersection of optimization methods, high-performance computing, and machine learning.

Scientific environment

The candidate will join the IRIT Laboratory (APO team) at ENSEEIHT (Toulouse-INP) as well as the international chair HAILSED at ANITI. The HAILSED chair, focusing on hybridizing AI and large-scale numerical simulations for engineering design, offers valuable opportunities to engage with experts in (scientific) machine learning, applied mathematics, scientific computing, numerical simulations, and high-performance computing. The candidate will also actively collaborate with the members of the DAIMOS project team, whose research interests span the development of exa-scale algorithms for large-scale meteo applications. This internship is intended to become a thesis as part of the DAIMOS project, which is a collaboration between IRIT, Inria Bordeaux, Sorbonne University, and Météo-France.

Candidate's profile

The ideal candidate should be motivated to continue in academic research with a PhD after this internship. They should have a strong interest in optimization, high-performance computing, machine learning, and domain-decomposition methods. A solid mathematical background combined with good programming skills will be essential for the successful completion of the project. Prior experience in one or more of the following areas would be highly advantageous:

- Programming in Python and familiarity with scientific computing libraries (e.g., `NumPy`, `SciPy`)
- Practical experience with machine learning frameworks such as `PyTorch`, `petsc4py`
- Practical experience with parallel computing and HPC computing platforms
- Understanding of numerical linear algebra and basics of domain decomposition (ASM/RAS)
- Knowledge of GNNs
- Knowledge of numerical optimization for training deep neural networks
- Good written and spoken English communication skills

The application

Interested candidates are required to submit an application that includes the following:

1. A comprehensive CV
2. A motivation letter detailing the applicant's research interests and reasons for applying

Please send your complete application in one single PDF file to Alena Kopaničáková (alena.kopanicakova@toulouse-inp.fr) and Julien Herrmann (julien.herrmann@irit.fr). The call is open until the position is filled.

Related literature

1. Serge Gratton, Alena Kopaničáková, and Philippe L. Toint. Multilevel objective-function-free optimization with an application to neural networks training. *SIAM Journal on Optimization*, 33(4):2772–2800, 2023

2. Serge Gratton, Alena Kopaničáková, and Philippe Toint. Recursive bound-constrained adagrad with applications to multilevel and domain decomposition minimization. *arXiv preprint arXiv:2507.11513*, 2025

3. Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S Yu. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1):4–24, 2020

4. Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1):61–80, 2008