Acceleration of Combinatorial Optimization Problem Solving Using Deep Learning Methods

6-month internship proposal

Laboratoire d'Informatique de l'X (LIX) - ORAILIX team
Palaiseau, France

Context

In industry, many stages of decision-making pipelines are formulated as combinatorial optimization problems, such as (Mixed) Integer Linear Programming (MILP/ILP) or Constrain Programming (CP) models. However, industrial instances are often very large and include numerous hard constraints that cannot be relaxed. As a result, these models involve a vast number of variables and constraints, making their resolution computationally expensive and, in many cases, infeasible within reasonable time limits.

Because of poor initialization, solvers may spend considerable time exploring suboptimal regions of the solution space before converging. A substantial body of research has therefore focused on constructing high-quality initial solutions to accelerate convergence. Traditionally, this has been achieved through heuristic methods that exploit domain knowledge about the problem [2]. More recently, deep learning–based warm-start techniques have emerged [5], where neural models generate promising initial solutions almost instantaneously, significantly speeding up the overall optimization process.

Objectives

Neural architectures can effectively learn structural patterns within families of combinatorial optimization problems (COPs), enabling the generation of useful initial solutions or partial variable assignment, reducing the effective decision space explored by solvers. As a prominent subclass of COPs, Mixed-Integer Linear Programs (MILPs) have attracted particular attention, with several works investigating such learning-based paradigms to accelerate their resolution when combined with traditional solvers [4, 3, 9].

This emerging line of research aims to establish a bridge between the formal language of optimization (variables, constraints, objectives) and the representational power of deep learning models. Graph Neural Networks (GNNs) and Transformers, in particular, have been leveraged to encode optimization instances as graph-based embeddings that capture the structural relationships between decision variables and constraints. This raises new research questions regarding how optimization instances can be more effectively represented in embedding space to enhance learning and inference performance. In the MILP setting, instances are often represented as bipartite graphs processed by GNN-based encoders. Investigating alternative architectures, such as Graph Transformers, could further improve solution quality and convergence speed [11]. Beyond graph-based models, graph-agnostic architectures inspired by sequence modeling have also been explored to mimic and enhance expert reasoning or solver trajectories [1, 7].

Alternatively, recent approaches based on discrete diffusion have emerged, enabling iterative refinement of solutions for COPs [13, 8]. These methods generate or improve candidate solutions through sequential denoising, offering a promising framework for refining variable assignments or heuristic strategies.

Extending existing machine learning–based warm-start techniques to other modeling paradigms beyond MILP, such as Constraint Programming (CP), represents a valuable research direction. Despite their expressiveness and suitability for a variety of industrial problems, CP formulations remain underexplored in the ML4CO literature [14].

The expected outcomes of this internship are threefold:

i. A literature review of recent advances in machine learning approaches for generating effective warm starts in combinatorial optimization solvers

ii. The benchmarking of promising methods on ILP standard datasets [4, 3, 9]

iii. The development of a deeper understanding of current limitations through experimentation on new instances of COPs representative of complex industrial problems, particularly routing and scheduling tasks such as the Capacitated Vehicle Routing Problem (CVRP) [10], the Multi-Agent Path Finding (MAPF) [12] and/or the Resource-Constrained Project Scheduling Problem (RCPSP) [6]. Insights gained from these analyses are expected to guide the design of a novel, higher-performing method developed during the internship.

If the results prove conclusive, the internship can lead to the submission of a scientific publication to a major machine learning conference or journal.

Candidate Profile

- Final-year student in "Grande Ecole d'Ingenieur" or Master's program.
- Strong interest in Deep Learning
- Keen interest in applied research
- Knowledge of Operations Research is a plus
- Experience with PyTorch in previous projects

How to Apply

The internship will be supervised by the ORAILIX research team of the Laboratoire d'Informatique de l'École Polytechnique (LIX) in Palaiseau. Interested candidates are invited to send a CV and a short motivation statement to one of the following people.

- Mathis LE BAIL, PhD Student (mathis.le-bail@polytechnique.edu)
- Clement ELLIKER, PhD Student (clement.elliker@polytechnique.edu)
- Sonia VANIER, HDR (sonia.vanier@polytechnique.edu)

References

[1] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Michael Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: reinforcement learning via sequence modeling. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, NIPS '21, Red Hook, NY, USA, 2021. Curran Associates Inc.

[2] Matteo Fischetti, Fred Glover, and Andrea Lodi. The feasibility pump. *Mathematical Programming*, 104(1):91–104, 2005.

[3] Zijie Geng, Jie Wang, Xijun Li, Fangzhou Zhu, Jianye HAO, Bin Li, and Feng Wu. Differentiable integer linear programming. In *The Thirteenth International Conference on Learning Representations*, 2025.

[4] Qingyu Han, L. Yang, Qian Chen, Xiang Zhou, Dong Zhang, Akang Wang, Ruoyu Sun, and Xiaodong Luo. A gnn-guided predict-and-search framework for mixed-integer linear programming. In *International Conference on Learning Representations*, 2023.

[5] Asunción Jiménez-Cordero, Juan Miguel Morales, and Salvador Pineda. Warm-starting constraint generation for mixed-integer optimization: A machine learning approach, 2022.

[6] Rainer Kolisch and Arno Sprecher. Psplib - a project scheduling problem library: Or software - orsep operations research software exchange program. *European Journal of Operational Research*, 96(1):205–216, 1997.

[7] Lucas Lehnert, Sainbayar Sukhbaatar, Paul Mcvay, Michael Rabbat, and Yuandong Tian. Beyond a*: Better planning with transformers via search dynamics bootstrapping. *ArXiv*, abs/2402.14083, 2024.

[8] Yang Li, Jinpei Guo, Runzhong Wang, and Junchi Yan. T2t: from distribution learning in training to gradient search in testing for combinatorial optimization. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS '23, Red Hook, NY, USA, 2023. Curran Associates Inc.

[9] Haoyang Liu, Jie Wang, Zijie Geng, Xijun Li, Yuxuan Zong, Fangzhou Zhu, Jianye HAO, and Feng Wu. Apollo-MILP: An alternating prediction-correction neural solving framework for mixed-integer linear programming. In *The Thirteenth International Conference on Learning Representations*, 2025.

[10] Denis Naddef and Giovanni Rinaldi. *Branch-And-Cut Algorithms for the Capacitated VRP*, pages 53–84.

[11] Ladislav Rampášek, Mikhail Galkin, Vijay Prakash Dwivedi, Anh Tuan Luu, Guy Wolf, and Dominique Beaini. Recipe for a general, powerful, scalable graph transformer. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS '22, Red Hook, NY, USA, 2022. Curran Associates Inc.

[12] Roni Stern. *Multi-Agent Path Finding – An Overview*, pages 96–115. 10 2019.

[13] Zhiqing Sun and Yiming Yang. Difusco: graph-based diffusion solvers for combinatorial optimization. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS '23, Red Hook, NY, USA, 2023. Curran Associates Inc.

[14] Yudong W. Xu, Wenhao Li, Scott Sanner, and Elias B. Khalil. Self-supervised transformers as iterative solution improvers for constraint satisfaction, 2025.