Boosting discrete optimization bounds using Deep Learning

Host team: Statistiques et Algorithmique pour la Biologie (SaAB)
Host laboratory:: MIA Toulouse, INRAE
Place: Castanet-Tolosan (near Toulouse), France
Encadrants: T. Schiex (thomas.schiex@inrae.fr), George Katsirelos (gkatsi@gmail.com), Simon de Givry (simon.de-givry@inrae.fr)
Renumeration: ≈ 550 € / mo

Topic: This Master's internship aims to explore the possibility of combining discrete optimization technology and deep learning, by using deep learning to accelerate the computation and improve the lower bounds used in a Branch and Bound discrete minimization process, to speed up search. This idea has already been explored for integer linear programming (Abbas and Swoboda 2024) and for solving the traveling salesman problem (Parjadis et al. 2024). In this internship, we want to explore it in the context of Cost Function Networks (CFN, (Cooper et al. 2010)). To learn bounding strategies, Deep Learning requires data, and we have a large set of problems at our disposal. The internship will be more specifically focused on Computational Protein Design instances (a problem with applications in health and green chemistry).

Context: For the discrete optimization side, we will rely on toulbar2 (see github.com/toulbar2/toulbar2 and toulbar2.github.io/toulbar2) as the discrete optimization tool. For exact discrete minimization, toulbar2 relies on a sophisticated Hybrid Best First branch & bound algorithm that relies on bounds that can be interpreted as Lagrangian or linear dual bounds of a linear continuous relaxation of the discrete CFN minimization problem (Cooper et al. 2010). The maximization of these dual bounds requires an optimization of the underlying Lagrangian multipliers. However, using linear programming or subgradient optimization is too expensive in practice. So, instead of exact optimization of these bounds, toulbar2 relies on fast approximations called “Soft Arc Consistencies” (Cooper et al. 2010).

The internship aims to build a deep learning architecture, likely a Graph Neural Network (GNN, (Corso et al. 2024)), that predicts the values of these multipliers from features of the instance being solved. The soft arc consistency heuristics may get stuck in local minima. The deep learning architecture could also help these heuristics get out of these minima to reach a global minimum more frequently. We will use PyTorch and likely rely on the geometric torch package for GNN here.

Beyond the definition of neural network architectures, the internship will require coding in Python, interfacing PyTorch with C++ code, and also running experimentations and training on GPUs. The internship may extend to a PhD thesis at the Artificial and Natural Intelligence Institute of Toulouse (ANITI).

References

Abbas, Ahmed, and Paul Swoboda. 2024. “Paul Swoboda: DOGE-Train: Discrete Optimization on GPU with End-to-End Training.” *AAAI*, 20623–31.

Cooper, M. C., S. de Givry, M. Sanchez, T. Schiex, M. Zytnicki, and T. Werner. 2010. “Soft Arc Consistency Revisited.” *Artificial Intelligence* 174 (7-8): 449–78.

Corso, Gabriele, Hannes Stark, Stefanie Jegelka, Tommi Jaakkola, and Regina Barzilay. 2024. “Graph Neural Networks.” *Nature Reviews. Methods Primers* 4 (1). https://doi.org/10.1038/s43586-024-00294-7.

Parjadis, Augustin, Quentin Cappart, Bistra Dilkina, Aaron Ferber, and Louis-Martin Rousseau. 2024. “Learning Lagrangian Multipliers for the Travelling Salesman Problem.” Schloss Dagstuhl - Leibniz-Zentrum für Informatik. https://doi.org/10.4230/LIPICS.CP.2024.22.