Master's Level Internship: Diffusion Models for Combinatorial Optimisation with Lagrangian Guidance

Advisors: Joseph Le Roux (leroux@lipn.fr) and Mathieu Lacroix (lacroix@lipn.fr)
Location: LIPN, Université Sorbonne Paris Nord – CNRS
Duration: 6 months

Context

Recent Machine Learning models for Combinatorial Optimisation such as (Sun & Yang, 2023) leverage denoising diffusion models with guidance (Ho & Salimans, 2021) to produce candidate solutions over combinatorial search spaces that take into account feasibility (Li et al., 2025). These models have shown tremendous success, as they exhibit state-of-the-art performance while enjoying short running times.

Diffusion models in the broad sense (including flows (Campbell et al., 2024) and flow matchings (Lipman et al., 2024) for instance) are generative probabilistic models that learn to generate samples from complex distributions by reversing a noising process that iteratively transforms data samples into uniformly distributed noise. They are used in combinatorial optimisation to design heuristics that cast MILPs as energy-based models (LeCun et al., 2006). They are useful to bypass traditional maximum likelihood estimation approaches where sampling is generally intractable.

Guidance is a generic term refering to methods designed to help the denoising process by incorporating knowledge at inference time (as opposed to training). While the best known of guidance is the classifier-free guidance (Ho & Salimans, 2021) where the diffusion process can be steered to generate samples conditioned on a specific class (*i.e.* generate an image conditioned on a specific class such as dog, cat etc. . .) recent diffusion models for Combinatorial Optimisation such as (Li et al., 2025) use guidance to inject information about constraint violations and reinforce objective optimisation.

Feasibility is modelled by a measure of the constraint violation. This measure is used in a similar fashion to Lagrangian Relaxation (Beasley, 1993; Frangioni, 2005) in order to reweight parts of the objective in order to penalize candidate solutions that violate constraints. Contrarily to Lagrangian Relaxation where reweighting is dynamic and calibrated for each instance specifically throught an iterative algorithms, the current approaches of guidance by violation treat penalty weights as hyper-parameters. This means that they must be carefully chosen, and that a value cherry-picked for a specific instance might not work for another.

Proposition

We propose to leverage the abundant literature on Lagrangian Relaxation to define a model where penalisation weights correspond to Lagrange multipliers and can be tuned to find the best compromise between objective optimisation and feasibility guarantees. The strength of the steering must be calibrated to enforce feasibility, as is the case for current approaches for guidance such as (Li et al., 2025), but also to reach a tight bound on the objective value. Starting from our methods (Demelas et al., 2025; Demelas et al., 2024) and the recent (Sander et al., 2025; Wiseman & Kim, 2019) we want to design a dynamic guidance that steer the prediction towards feasible solution when needed and only when needed.

A possible approach is to train two diffusion models: one in charge of predicting solutions, the other one in charge of predicting the penalty. The interplay between these two approaches will be the subject of this internship.

Application

Candidates are expected to be graduating at Master level with a strong background in Machine Learning. Candidates are not required to be knowledgeable in Combinatorial Optimization but a background in this topic will be appreciated.

This is a 6 month internship, to apply please send email with a CV, a cover letter and a transcript stating clearly your period of availability.

References

Beasley, J. E. (1993). Lagrangian relaxation. In *Modern heuristic techniques for combinatorial problems* (pp. 243–303). John Wiley & Sons, Inc.

Campbell, A., Yim, J., Barzilay, R., Rainforth, T., & Jaakkola, T. (2024, 21–27 Jul). Generative flows on discrete state-spaces: Enabling multimodal flows with applications to protein co-design. In R. Salakhutdinov, Z. Kolter, K. Heller, A. Weller, N. Oliver, J. Scarlett, & F. Berkenkamp (Eds.), *Proceedings of the 41st international conference on machine learning* (pp. 5453–5512, Vol. 235). PMLR. https://proceedings.mlr.press/v235/campbell24a.html

Demelas, F., Roux, J. L., Frangioni, A., Lacroix, M., Traversi, E., & Calvo, R. W. (2025). Bundle network: A machine learning-based bundle method. https://arxiv.org/abs/2509.24736

Demelas, F., Le Roux, J., Lacroix, M., & Parmentier, A. (2024). Predicting lagrangian multipliers for mixed integer linear programs. *Forty-first International Conference on Machine Learning*. https://openreview.net/forum?id=aZnZOqUOHq

Frangioni, A. (2005). About Lagrangian Methods in Integer Optimization. *Annals of Operations Research*, *139*(1). http://link.springer.com/10.1007/s10479-005-3447-9

Ho, J., & Salimans, T. (2021). Classifier-free diffusion guidance. *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*. https://openreview.net/forum?id=qw8AKxfYbI

LeCun, Y., Chopra, S., Hadsell, R., Ranzato, M., & Huang, F. (2006). A tutorial on energy-based learning. *Predicting structured data*, *1*(0).

Li, H., Yuan, H., Zhang, H., Lin, J., Ge, D., Wang, M., & Ye, Y. (2025). Fmip: Joint continuous-integer flow for mixed-integer linear programming. https://arxiv.org/abs/2507.23390

Lipman, Y., Havasi, M., Holderrieth, P., Shaul, N., Le, M., Karrer, B., Chen, R. T. Q., Lopez-Paz, D., Ben-Hamu, H., & Gat, I. (2024). Flow matching guide and code. https://arxiv.org/abs/2412.06264

Sander, M. E., Roulet, V., Liu, T., & Blondel, M. (2025). Joint learning of energy-based models and their partition function. *CoRR*. http://arxiv.org/abs/2501.18528v1

Sun, Z., & Yang, Y. (2023). Difusco: Graph-based diffusion solvers for combinatorial optimization. *CoRR*. http://arxiv.org/abs/2302.08224v2

Wiseman, S., & Kim, Y. (2019). Amortized bethe free energy minimization for learning mrfs. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, & R. Garnett (Eds.), *Advances in neural information processing systems* (Vol. 32). Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2019/file/dc554706afe4c72a60a25314cbaece80-Paper.pdf