M2 internship and PhD research project

Machine learning based hybrid strategies for combinatorial optimization

Keywords: combinatorial optimization, vehicle routing problem with time windows, machine learning, structured learning, implicit layers, statistical learning.

1 Context

Rungis Market is the largest fresh produce market in the world, in particular, it supplies 63% of the food consumed in Île-de-France. The company Califrais is its official digital and logistics operator through the platform rungismarket.com. With the aim of digitalizing and decarbonizing the food supply chain, Califrais develops logistics solutions based on algorithms at the crossroads of machine learning, statistics and optimization.

Califrais has strong connections with public research and works with academic partners such as the Laboratoire de Probabilités, Statistique et Modélisation (LPSM), Sorbonne Université and the CNRS. This privileged context has led to numerous publications in leading conferences of machine learning, such as ICML and NeurIPS.

We are offering this research internship as part of a collaboration between Califrais and the LPSM. Its aim is to study recent tools at the intersection of machine learning and combinatorial optimization that allow to jointly "learn and optimize". It is funded by the French Agency for Ecological Transition (ADEME) under the call "Logistique 4.0" won by Califrais and Sorbonne Université. The internship is designed to lead to a Ph.D. in October 2025.

2 Subject

Califrais has a history of several years of daily order and delivery data for thousands of fresh products and hundreds of customers. Extracting information from this large amount of data is critical for a variety of tasks: predictive analysis of demand or customer satisfaction, inventory optimization of perishable goods or routing optimization. However, these data present many challenges and domain-specific characteristics; in particular, they have a lot of noise and non-stationarity. This is due to the specificity of fresh products: many external factors affect both supply and demand of these products, such as weather, shortages, political and global economic factors,...

Several research topics at Califrais fall into the domain of discrete optimization but parameters of these optimization problems are often unknown quantities. For example, to optimize an inventory, it is necessary to take into account future demands which are unknown. Similarly, in routing problems, travel times are unknown but these are crucial quantities in the optimization to satisfy customer delivery constraints, modeled as time window in which the delivery must take place. To handle these unknown quantities, the traditional procedure is to have a two-steps approach: first predict the unknown quantities, and then optimize. However, this approach is not completely satisfactory since the predicted quantities are then treated by the optimization algorithm as true quantities. In other words, their uncertainty is not accounted for. When dealing with noisy signals, which is the case for food supply chain, this leads to optimization algorithms that "overfit" on prediction errors.

Therefore, it is natural to turn to approaches that tackle both learning and optimization tasks together. This is a rich and diverse domain; see, for example, Donti et al. (2017); Bengio et al. (2021); Bai et al. (2023); Mandi et al. (2024); Sadana et al. (2025); Vivier-Ardisson et al. (2024). A recent and promising direction is the use of combinatorial optimization layers together with machine learning models (Berthet et al., 2020; Dalle et al., 2022). The idea is to treat a combinatorial optimization solver as a black-box component in a machine learning pipeline, through wich gradients can be back-propagated. The pipeline can be summarized as follows:

$$\text{Input } x \in \mathcal{X} \to \text{ ML prediction } \to \text{ Parameters } \theta = h_w(x) \to \text{ CO layer } \to \text{ Solution } y = f(\theta). \tag{1}$$

An input x is given to a learning algorithm h_w, where w are the parameters of the model, which outputs a vector $\theta = h_w(x)$. This parameter θ defines the optimization objective, which typically writes as

$$\underset{v \in \mathcal{Y}}{\arg\max}\ \theta^\top v, \tag{2}$$

where $\mathcal{Y}$ is a discrete set. Some algorithm, either an exact solver or a heuristics, outputs a solution of (2), denoted by $y = f(\theta)$. The goal is then to find the parameters w that will give good final solutions y. Training is typically done in a supervised setting: given a dataset of pairs of problems and solutions $\{(x_i, y_i)\}_{i=1,\dots,n}$ and a distance function ℓ on the space $\mathcal{Y}$, the goal is to minimize the loss

$$\mathcal{L}_n = \sum_{i=1}^{n} \ell\big(f(h_w(x_i)), y_i\big).$$

The first objective of the internship will be to apply this general framework to the routing problem of Califrais. A starting point can be the work by Baty et al. (2024), which explores a related dynamic vehicle problem. Our problem is the following: given a graph, denoted by $(\mathcal{V}, \mathcal{A})$, where $\mathcal{V}$ denotes the set of vertices ($|\mathcal{V}| = n$) and $\mathcal{A}$ the set of edges, the goal is to partition this graph in different routes, each corresponding to a vehicle. One of the vertices of the graph corresponds to the depot (think, Rungis) from which all the routes have to begin and come back, and every other vertex corresponds to a client. Each vertex is associated to a target delivery time window $[a_i, b_i]$. The objective is to minimize the travel time under the constraint that the delivery at node i has to happen within the time window. The problem is that the time to go from one vertex to another is unknown a priori, which is where machine learning tools come into play. Formally, if we denote by $\mathcal{Y} \subset \{0,1\}^{n \times n}$ the set of feasible solutions, where a solution $y \in \mathcal{Y}$ is represented as: $y_{ij} = 1$ if a route goes from node i to node j, then, theoretically, we would like to minimize

$$\min_{y \in \mathcal{Y}} \sum_{(i,j) \in \mathcal{A}} c_{i,j}(y) y_{i,j}, \tag{3}$$

where $c_{i,j}(y)$ denotes the travel time to go from node i to node j for the solution y. The complicated aspect is that the travel time depends on the hour of the day, meaning that it depends on *all the other vertices visited before node* j on the route, and not only on the previously visited node i. In other words, in (3), we do not have a cost $c_{i,j}$ but a cost $c_{i,j}(y)$. Following the pipeline presented above, a promising direction would be to replace these complicated costs $c_{i,j}(y)$ by a ML algorithm, meaning that problem (3) is replaced by

$$\min_{y \in \mathcal{Y}} \sum_{(i,j) \in \mathcal{A}} (h_w(x))_{i,j} y_{i,j}. \tag{4}$$

A second step would be to transform this vehicle routing problem into a dynamic one, splitting the delivery period into several epochs, and to assume the travel times as constant within each epoch.

Having familiarised with these tools and implemented an initial solution for the vehicle routing application, we will take a step back and look at the theoretical guarantees associated with pipelines of the form (1). Several directions can be undertaken, such as combining statistical learning guarantees on the ML prediction algorithm with results on the combinatorial optimization side, or investigating other learning paradigms such as reinforcement learning or active learning.

3 Practical conditions

This internship is part of a collaboration between Califrais and the LPSM at Sorbonne Université.

- Location: the internship will take place in both locations (Paris 75005 and Paris 75010).
- Duration: 5 to 6 months between February and September 2025.
- Profile: M2 research in applied mathematics or third year of engineering school.
- Grant: $4,35$€/h, being around 600€/month.

The supervising team will consist of:

- Adeline Fermanian, Califrais
- Maxime Sangnier, Sorbonne Université

4 Application

The application file containing:

- a resume;
- a cover letter;
- a transcript of grades (Bachelor's and Master's degrees);

should be sent by mail to adeline.fermanian@califrais.fr and maxime.sangnier@sorbonne-universite.fr.

References

R. Bai, X. Chen, Z.-L. Chen, T. Cui, S. Gong, W. He, X. Jiang, H. Jin, J. Jin, G. Kendall, et al. Analytics and machine learning in vehicle routing research. *International Journal of Production Research*, 61(1): 4–30, 2023.

L. Baty, K. Jungel, P. S. Klein, A. Parmentier, and M. Schiffer. Combinatorial optimization-enriched machine learning to solve the dynamic vehicle routing problem with time windows. *Transportation Science*, 2024.

Y. Bengio, A. Lodi, and A. Prouvost. Machine learning for combinatorial optimization: A methodological tour d'horizon. *European Journal of Operational Research*, 290(2):405–421, 2021.

Q. Berthet, M. Blondel, O. Teboul, M. Cuturi, J.-P. Vert, and F. Bach. Learning with differentiable pertubed optimizers. *Advances in neural information processing systems*, 33:9508–9519, 2020.

G. Dalle, L. Baty, L. Bouvier, and A. Parmentier. Learning with combinatorial optimization layers: a probabilistic approach. *arXiv preprint arXiv:2207.13513*, 2022.

P. L. Donti, B. Amos, and J. Z. Kolter. Task-based end-to-end model learning. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

J. Mandi, J. Kotary, S. Berden, M. Mulamba, V. Bucarey, T. Guns, and F. Fioretto. Decision-focused learning: Foundations, state of the art,benchmark and future opportunities. *Journal of Artificial Intelligence Research*, 81:1623–1701, 2024.

U. Sadana, A. Chenreddy, E. Delage, A. Forel, E. Frejinger, and T. Vidal. A survey of contextual optimization methods for decision-making under uncertainty. *European Journal of Operational Research*, 320(2):271–289, 2025.

G. Vivier-Ardisson, A. Forel, A. Parmentier, and T. Vidal. Cf-opt: Counterfactual explanations for structured prediction. In *Proceedings of the 41 st International Conference on Machine Learning*, volume 235. PMLR, 2024.