



HAL
open science

Learning structured approximations of combinatorial optimization problems.

Axel Parmentier

► **To cite this version:**

Axel Parmentier. Learning structured approximations of combinatorial optimization problems.. 2023.
hal-03281894v2

HAL Id: hal-03281894

<https://hal.science/hal-03281894v2>

Preprint submitted on 27 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Learning structured approximations of combinatorial optimization problems.

Axel Parmentier
CERMICS, Ecole des Ponts, Marne-la-Vallée, France
axel.parmentier@enpc.fr

January 27, 2023

Abstract

Machine learning pipelines that include a combinatorial optimization layer can give surprisingly efficient heuristics for difficult combinatorial optimization problems. Three questions remain open: which architecture should be used, how should the parameters of the machine learning model be learned, and what performance guarantees can we expect from the resulting algorithms? Following the intuitions of geometric deep learning, we explain why equivariant layers should be used when designing such pipelines, and illustrate how to build such layers on routing, scheduling, and network design applications. We introduce a learning approach that enables to learn such pipelines when the training set contains only instances of the difficult optimization problem and not their optimal solutions, and show its numerical performance on our three applications. Finally, using tools from statistical learning theory, we prove a theorem showing the convergence speed of the estimator. As a corollary, we obtain that, if an approximation algorithm can be encoded by the pipeline for some parametrization, then the learned pipeline will retain the approximation ratio guarantee. On our network design problem, our machine learning pipeline has the approximation ratio guarantee of the best approximation algorithm known and the numerical efficiency of the best heuristic.

1 Introduction

In the last few years, more and more attention have been given to the construction of machine learning algorithms which, given an input x , can predict an output y in a combinatorially large set $\mathcal{Y}(x)$. An approach that is getting more popular to address this problem consists in embedding a combinatorial optimization (CO) layer in a machine learning pipeline. As illustrated on Figure 1, the resulting pipeline typically chains a statistical model, a combinatorial optimization problem, and possibly a post-processing algorithm.




Figure 1: A machine learning pipeline with a combinatorial optimization (CO) layer

Such pipelines can be used as heuristics for difficult combinatorial optimization problems. Let us consider a combinatorial optimization problem of interest

$$\min_{z \in \mathcal{Z}(x)} C(z, x). \quad (\text{Pb})$$

Here x is an instance in a *set of instance* \mathcal{X} , and $\mathcal{Z}(x)$ denotes the set of feasible solutions of x . Contrary to what is usual in combinatorial optimization, we include the instance x in the objective function $c(z; x)$.

When we use our machine learning pipeline to solve (Pb), we use the statistical model f_w to obtain the parameter θ of the auxiliary combinatorial optimization problem

$$\min_{y \in \mathcal{Y}(x)} g(y, \theta), \quad (\text{CO-layer})$$

and then decode the solution y of this problem into a solution $z \in \mathcal{Z}(x)$ of the initial problem. Such a pipeline is useful when we have much more efficient algorithms for our combinatorial optimization layer problem (CO-layer) than for the problem of interest (Pb).

Since the preprocessing h is assumed deterministic, it will not play a major role on the learning algorithm. It will therefore be convenient to bring back the cost on $\mathcal{Y}(x)$. For y in $\mathcal{Y}(x)$, we define

$$c(y, x) = C(h(y), x).$$

Running example: Two stage spanning tree Let $G = (V, E)$ be an undirected graph, and S be a finite set of scenarios. The objective is to build a spanning tree on G of maximum cost on a two stage horizon. Building edge e in the first stage costs $c_e \in \mathbb{E}$, while building it in the second stage under scenario s costs $d_{es} \in \mathbb{R}$. The decision maker does not know the scenario s when it chooses which first stage edges to build. Denoting \mathcal{T} the set of spanning trees, we can formulate the problem as

$$\min \left\{ \sum_{e \in E_1} c_e + \frac{1}{|S|} \sum_{e \in E_s} d_{es} : E_1 \cap E_s = \emptyset \text{ and } (V, E_1 \cup E_s) \in \mathcal{T} \text{ for all } s \text{ in } S \right\}. \quad (1)$$

When we restrict ourselves to $c_e \leq 0$ and $d_{es} \leq 0$, we obtain the two stage maximum weight spanning tree. [Escoffier et al. \[2010\]](#) show that this restriction is APX-complete, and introduce a 2-approximation algorithm for the maximization problem, which translates into a 1/2-approximation algorithm for the minimization problem.

Running example pipeline Remark that an optimal solution of the single scenario version of the problem

$$\min \left\{ \sum_{e \in E_1} \bar{c}_e + \frac{1}{|S|} \sum_{e \in E_2} \bar{d}_e : E_1 \cap E_2 = \emptyset \text{ and } (V, E_1 \cup E_2) \in \mathcal{T} \right\} \quad (2)$$

is a minimum weight spanning tree on G with edge weights $\min(\bar{c}_e, \bar{d}_e)$. It can therefore be easily solved using Kruskal’s algorithm, and we therefore suggest using (2) as combinatorial optimization layer (**CO-layer**). Hence, we have $\theta = (c_e, d_e)_{e \in E}$ and $\Theta(x) = \mathbb{R}^{2E}$.

Our decoder h rebuilds a solution z of (1) from a solution $y = (\bar{E}_1, \bar{E}_2)$ of (2). It relies on the following result. Given a forest F , Kruskal’s algorithm can be adapted to find a minimum weight spanning tree containing F . We take \bar{E}_1 as the first stage solution of (1), and use the variant of Kruskal’s algorithm with edge weights d_s to rebuild the E_s . We then compare this solution to the optimal solution $E_1 = \emptyset$ and return the best of the two as z .

Structure of the combinatorial optimization layer. When building a solution pipeline for a combinatorial algorithm, we typically want our pipeline to be able to address instances of very different size: Instances of our running example may have 20 or 1000 edges. It means that the graph G used in the combinatorial optimization layer (2) depends on the instance x of (1), and hence the parameter θ belongs to the set \mathbb{R}^{2E} which also depends on x . This is the reason why, in our pipeline, the set of solutions $\mathcal{Y}(x)$ and the parameter space $\Theta(x)$ both depend on x . On the contrary, since we want to use the same model and hence the same f_w on different instances, the space \mathcal{W} does not depend on x . This raises the question of *how to build statistical model f_w whose output dimension depends on the input dimension*. More generally, such an approach can work only if (**CO-layer**) retains most of the “structure” of (Pb).

Learning algorithm. Finally, the purpose of the learning algorithm is to find a parameter $w \in \mathcal{W}$ such that the pipeline outputs a good solution of (Pb). Approaches in the literature typically use a *learning by imitation approach*, with a training set $(x_1, z_1), \dots, (x_n, z_n)$ containing instances of (Pb) and their hard problem solution. A drawback of such an approach is that it requires another solution algorithm for (Pb) to compute the z_i . In this paper, we focus on the *learning by experience* setting where the training set contains only instances x_1, \dots, x_n .

Related works. The interactions between combinatorial optimization and machine learning is an active research area [Bengio et al., 2021]. Combinatorial optimization layers in deep learning belong to the subarea of end-to-end learning methods for combinatorial optimization problems recently surveyed by Kotary et al. [2021]. This field can be broadly classified in two subfields. Machine learning augmented combinatorial optimization uses machine learning to take heuristic decisions within combinatorial optimization algorithms. We survey here combinatorial optimization augmented machine learning, which inserts combinatorial optimization oracles within machine learning pipelines.

Structured learning approaches were the first to introduce these methods in the early 2000s [Nowozin, 2010] in the machine learning community. They mainly considered maximum a posteriori problems in probabilistic graphical models as combinatorial optimization layers, with applications to computer vision, and sorting algorithms with applications to ranking. They were generally trained using the structured Hinge loss or a maximum likelihood estimator. A renewed interest for optimization layers in deep learning pipeline has emerged in the last few years has emerged in the machine learning community, and notably continuous optimization layers [Amos and Kolter, Blondel et al., 2022]. Remark that these pipelines are generally trained using a learning by imitation paradigm.

We focus here on combinatorial optimization layers. Among these, linear optimization layers have received the most attention. Two challenges must be addressed. First, since the mapping that associated to the objective parameter vector θ the output y is piecewise constant, and deep learning networks are generally trained using stochastic gradient descent, meaningful approximations of the must be proposed gradient [Vlastelica et al.]. Second a loss quantifying the error between its target must be proposed. Blondel et al. address these challenges with an elegant solution based on convex duality: the linear objective is regularized with a convex penalization, which leads to meaningful gradients. Fenchel Young inequality in convex duality then gives a natural definition of the loss function. Berthet et al. [2020] have shown that this approach can be extended to the case where a random perturbation is added to the objective instead of a convex regularization. When it comes to integer linear programs, Mandi et al. [2020] suggest using the linear relaxation during the learning phase.

The author recently introduced the idea of building heuristics for hard combinatorial optimization problems with pipelines with combinatorial optimization layers [Parmentier, 2021]. The closest contribution to our learning by experience setting is the smart predict then optimize method of Elmachtoub and Grigas [2021]. It considers the case where there is no decoder and the cost function $c(y, x) = \theta^*x$ is actually the linear objective of the combinatorial optimization layer $g(y, \theta) = \theta y$ for an unknown true parameter θ^* . They propose a generalization of the structured Hinge loss to that setting.

However, to the best of our knowledge, two aspects of pipelines with combinatorial optimization layers have not been considered in the literature. First, the general learning by experience setting where only instances of the hard optimization problems are available has not been considered. Second, there is no guarantee on the quality of the solution returned by the pipeline. The purpose of this paper is to address these two issues.

Contributions We make the following contributions.

1. The design of the learning pipelines, and notably the choice of (CO-layer) and f_w is critical for the performance of the resulting algorithm. We illustrate on three applications among which our running example how to build such pipelines.
2. A natural way of formulating the learning problem consists in minimizing the loss

defined average cost of the solution z_i returned by our pipeline for instance x_i

$$\frac{1}{n} \sum_{i=1}^n c \left(\arg \min_{y \in \mathcal{Y}(x_i)} g(y, f_{\mathbf{w}}(x_i)), x_i \right)$$

We introduce a regularized version of this loss. And we show with extensive numerical experiments that, despite the non-convexity of this loss, when the dimension of \mathcal{W} is moderate, i.e., non-greater than 100, solving this problem with a global black-box solver leads to surprisingly efficient pipelines.

3. Leveraging tools from statistical learning theory, we prove the convergence of the learning algorithm toward the approximation with the best expected loss, and an upper bound on the convergence speed.
4. We deduce from these statistical learning results that, under some hypotheses on the pipeline, the learned algorithm is an approximation algorithm for (Pb). These hypotheses are notably satisfied by our solution pipeline for (1).

Remark that, in this paper, we do not try to approximate difficult constraint. We only try to approximate difficult objectives. The paper is organized as follows. Section 2 introduce two additional examples and explain how to build pipelines. Section 3 formulates the learning by experience problem and introduces algorithms. Section 4 introduces the convergence results and the approximation ratio guarantee. Finally, Section 5 details the numerical experiments.

2 Designing pipelines with combinatorial optimization layers

In this section, we give a methodology to build pipelines with combinatorial optimization layers. We illustrate it on our running example and on two applications previously introduced by the author. We start with the description of these applications, which follows the papers which introduced them [Parmentier, 2021, Parmentier and T'Kindt, 2021].

2.1 Stochastic vehicle scheduling problem.

Stochastic vehicle scheduling problem Let V be a set of tasks that should be operated using vehicles. For each task v in V , we suppose to have a scheduled start time t_v^b in \mathbb{Z}_+ and a scheduled end time t_v^e in \mathbb{Z}_+ . We suppose $t_v^e > t_v^b$ for each task v in V . For each pair of tasks (u, v) , the travel time to reach task v from task u is denoted by $t_{(u,v)}^{\text{tr}}$. Task v can be operated after task u using the same vehicle if

$$t_v^b \geq t_u^e + t_{(u,v)}^{\text{tr}}. \tag{3}$$

We introduce the digraph $D = (V, A)$ with vertex set $V = T \cup \{o, d\}$ where o and d are artificial origin and destination vertices. The arc set A contains the pair (u, v) in T^2 if v can be scheduled after task u , as well as the pairs (o, v) and (v, d) for all v in V . An o - d path P represents a sequence of tasks operated by a vehicle. A feasible solution is a partition of V into o - d paths. If we denote by $c(P, x)$ the cost of operating the sequence corresponding to the o - d path P , and by \mathcal{P}_{od} the set of o - d paths, the problem can be modeled as follows.

$$\min_z \sum_{P \in \mathcal{P}_{od}} c(P; x) z_P, \quad (4a)$$

$$\text{s.t. } \sum_{P \ni v} z_P = 1, \quad \forall v \in V(x) \setminus \{o, d\}, \quad (4b)$$

$$z_P \in \{0, 1\}, \quad \forall P \in \mathcal{P}_{od}(x), \quad (4c)$$

Up to now, we have described a generic vehicle scheduling problem. Let us now define our stochastic vehicle scheduling problem by giving the definition of $c(P, x)$. Let Ω be a set of scenarios. For each task v , we have a random start time ξ_v^b and a random end time ξ_v^e , and for each arc (u, v) , we have a random travel time $\xi_{(u,v)}^{\text{tr}}$. Hence, $\xi_v^b(\omega)$, $\xi_v^e(\omega)$, and $\xi_{(u,v)}^{\text{tr}}(\omega)$ are respectively the beginning time of v , end time of v , and travel time between u and v under scenario ω in Ω . We define $\xi_o^e = 0$ and $\xi_d^b = +\infty$.

Given an o - v path P , we define recursively the end-time τ_P of P as follows.

$$\tau_P = \begin{cases} 0, & \text{if } P \text{ is the empty path in } o, \\ \xi_v^e + \max(\tau_Q + \xi_a^{\text{tr}} - \xi_v^b, 0), & \text{if } P = Q + a \text{ for some path } Q \text{ and arc } a. \end{cases} \quad (5)$$

Equation (5) models the fact that a task can be operated by a vehicle only when the vehicle has finished the previous task: The vehicle finishes Q at τ_Q , and arrives in v at $\tau_Q + \xi_a^{\text{tr}}$ with delay $\max(\tau_Q + \xi_a^{\text{tr}} - \xi_v^b, 0)$. The total delay Δ_P along a path P is therefore defined recursively by

$$\Delta_P = \begin{cases} 0, & \text{if } P \text{ is the empty path in } o, \\ \Delta_Q + \max(\tau_Q + \xi_a^{\text{tr}} - \xi_v^b, 0), & \text{if } P = Q + a \text{ for some path } Q \text{ and arc } a. \end{cases} \quad (6)$$

Finally, we define the cost of an o - d path P as

$$c(P; x) = c^{\text{veh}} + c^{\text{del}} \mathbb{E}(\tau_P) \quad (7)$$

where c^{veh} in \mathbb{Z}_+ is the cost of a vehicle and c^{del} in \mathbb{Z}_+ is the cost of a unit delay. Practically, we use a finite set of scenarios Ω , and compute the expectation as the average on this set.

CO layer: usual vehicle scheduling problem The usual vehicle scheduling problem can also be formulated as (4), the difference being that now the path can be decomposed as the sum of the arcs cost

$$\bar{c}_P = \sum_{a \in P} \bar{c}_a \quad \text{with } \bar{c}_a \in \mathbb{R}. \quad (8)$$

It can be reduced to a flow problem on D and efficiently solved using flow algorithms or linear programming. In Equation (8) and in the rest of the paper, we use an overline to denote quantities corresponding to the easy problem.

2.2 Single machine scheduling problem.

Scheduling problem $1|r_j|\sum_j C_j$. n jobs must be processed in a single machine. Jobs cannot be interrupted once launched. Each job has a processing time p_j , and a release time r_j in \mathbb{R} . That is, job j cannot be started before r_j , and once started, it takes p_j to complete it. A solution is a schedule $s = (j_1, \dots, j_n)$, i.e., a permutation of $[n]$ that gives the order in which jobs are processed. Using the convention $C_{j_0} = 0$, the completion time of jobs in s are defined as

$$C_{j_i} = \max(r_j, C_{j_{i-1}}) + p_{j_i}.$$

The objective is to find a solution minimizing $\sum_j C_j$. This problem is strongly NP-hard.

Combinatorial optimization layer: $1||\sum_j C_j$. The easy problem is obtained when there is no release time, and only jobs processing times \bar{p}_j . Jobs completion times are therefore given by

$$\bar{C}_{j_i} = \bar{C}_{j_{i-1}} + \bar{p}_{j_i}.$$

Again, we use an overline to denote quantities of the easy problem. An optimal schedule is obtained using the shortest processing time first (SPT) rule, that is, by sorting the jobs by increasing p_j .

2.3 Constructing pipelines

In this section, we explain how to build our learning pipelines.

Combinatorial Optimization layer and decoder. The choice of the combinatorial optimization layer and the decoder are rather applications dependent. Two practical aspects are important. First, we must have a practically efficient algorithm to solve (**CO-layer**). Second, it must be easy to turn solutions of (**CO-layer**) into solution of (**Pb**). That is, either the solutions of (**CO-layer**) and (**Pb**) coincide, or we must have a practically efficient algorithm h that turns a solution of (**CO-layer**) into a solution of (**Pb**).

Structure of x and generalized linear model. As we indicated in the introduction, a practical difficulty in the definition of our statistical model $f_{\mathbf{w}}$ is that the size of its output $\boldsymbol{\theta}$ in $\Theta(x)$ depends on the instance x . Unfortunately, statistical models generally output vectors of fixed size. Let us pinpoint a practical way of addressing this difficulty with a generalized linear model. Let $\mathcal{I}(x)$ be the *structure* x , i.e., the set of dimensions i of $\Theta(x)$. We suggest defining a feature mapping

$$\phi : (i, x) \mapsto \phi(i, x)$$

that associates to an instance and a dimension i in $\mathcal{I}(x)$ a feature vector $\phi(i, x)$ describing the main properties of i as a dimension of x . We then define

$$f_{\mathbf{w}} : x \mapsto \boldsymbol{\theta} \quad \text{with} \quad \boldsymbol{\theta} = (\theta_i)_{i \in \mathcal{I}(x)} \quad \text{and} \quad \langle \mathbf{w} | \phi(i, x) \rangle.$$

Feature description	$\phi((e, \mathbf{first}), x)$	$\phi((e, \mathbf{second}), x)$
First stage cost	c_e	0
Second stage average cost	0	$\sum_s d_{es}/ S $
Quantiles of second stage cost	0	$Q[(d_{es})_s]$
Quantiles of neighbors first stage cost	$Q[(c_{e'})_{e' \in \delta(u) \cup \delta(v)}]$	0
Quantiles of neighbors second stage cost	0	$Q[(d_{e's})_{e' \in \delta(u) \cup \delta(v), s \in S}]$
“Is edge in first stage MST ?”	$\mathbb{1}^{\text{MST}}(e, (c_e)_{e \in E})$	0
Quantiles of “Is edge in second stage MST quantile ?”	0	$Q\left[\left(\mathbb{1}^{\text{MST}}(e, (b_{es})_{e \in E})\right)_{s \in S}\right]$
Quantiles of “Is first stage edge in best stage MST quantile ?”	$Q\left[\left(\mathbb{1}^{\text{MST}}(e, (b_{es})_{e \in E})\right)_{s \in S}\right]$ and $c_e \leq d_{es}$	0
Quantiles of “Is second stage edge in best stage MST quantile ?”	0	$Q\left[\left(\mathbb{1}^{\text{MST}}(e, (b_{es})_{e \in E})\right)_{s \in S}\right]$ and $c_e > d_{es}$

Note: MST stands for Minimum Weight Spanning Tree, $b_{es} = \min(c_e, d_{es})$, $Q[\mathbf{a}]$ gives the quantiles of a vector \mathbf{a} seen as a sampled distribution, and $\mathbb{1}^{\text{MST}}(e, (\tilde{c}_e)_e)$ is equal to 1 if e is in the minimum spanning tree for edge weights $(\tilde{c}_e)_e$.

Table 1: Two stage spanning tree features of edge $e = (u, v)$.

In summary, $f_{\mathbf{w}}$ can output parameters $\boldsymbol{\theta}$ whose dimension depends on x because it applies the same predictor $(i, x) \mapsto \langle \mathbf{w} | \phi(i, x) \rangle$ to predict the value for θ_i for the different dimensions in $\mathcal{I}(x)$.

Illustration on our applications. For instance, let us consider our running example on *two stage spanning tree problem*. Given an instance x , we must define the first and second stage costs \bar{c}_e and \bar{d}_e for each edge $e \in E$. We can therefore define $\mathcal{I}(x)$ as $\{(e, \mathbf{stage}) : e \in E, \mathbf{stage} \in \{\mathbf{first}, \mathbf{second}\}\}$. The details of the features used is described in Table 1. For the stochastic vehicle scheduling problem, all we have to do is to define the arc costs \bar{c}_a . Hence, we can define $\mathcal{I}(x) = A$. And for the single machine scheduling problem, we only have to define the processing times \bar{p}_j . Hence, $\mathcal{I}(x) = \{1, \dots, n\}$.

Encoding information on an element as part of an instance Let us finally introduce two generic techniques to build interesting features. The features in Table 1 rely on these two techniques. The first technique enables to compare dimension i to the other ones in $\mathcal{I}(x)$. To that purpose, we define a statistic $\alpha : (i, \mathcal{X}) \mapsto \alpha(i, x)$, and considers $f(\boldsymbol{\rho}_e)$ as a realization of the random variable

$$\begin{aligned} \mathcal{A} : \mathcal{I}(x) &\rightarrow \mathbb{R} \\ i &\mapsto \alpha(i, x) \end{aligned}$$

and take some relevant statistics on the realization $\mathcal{A}(i)$ of \mathcal{A} , such as the value of the cumulative distribution function of F in $\alpha(i, x)$. For instance, when considering a job j of $1|r_j|\sum_j C_j$ with parameter (r_j, p_j) , if we define $\alpha(j, x) = r_j + p_j$, we obtain as feature the rank (divided by n) of feature j in the schedule where we sort the jobs by increasing $r_j + p_j$, a statistic known to be interesting and used in dispatching rules.

The second technique is to *explore the role of i in the solution of a very simple optimization problem*. A natural way of building features is to run a fast heuristic on the instance x and seek properties of i in the resulting solution. For instance, the preemptive version of $1|r_j|\sum_j C_j$, where jobs can be stopped, is easy to solve. Statistics such as the number of times job j is preempted in the optimal solution can be used as features.

Equivariant layers. A lesson from geometric deep learning [Bronstein et al., 2021] is that good neural network architecture should respect the symmetries of the problem. Let \mathcal{S} be a symmetry of the problem. A layer h in a neural network is said to be equivariant with respect to \mathcal{S} if $h(\mathcal{S}(x)) = \mathcal{S}(h(x))$. In our combinatorial optimization setting, there is one natural symmetry. The solution predicted y should not depend on the indexing of the variables using in the combinatorial optimization problem : Given a permutation of these variables in the instance x , the solution y should be the permuted solution. Combinatorial optimization layers are naturally equivariant with respect to this symmetry. The generalized linear model above is a simple example of equivariant layer.

3 Learning by experience

We now focus on how to learn pipelines with a combinatorial optimization layer. Given a training set composed of representative instances, the learning problem aims at finding a parameter \mathbf{w} such that the output $z(\mathbf{w})$ of our pipeline has a small cost.

As we mentioned in the introduction, the literature focuses on the learning by imitation setting. In that case, the training set $(x_1, y_1), \dots, (x_n, y_n)$ contains instances and target solution of the prediction problem (**CO-layer**), the learning problem can be formulated as

$$\min \frac{1}{n} \sum_{i=1}^n \ell(\boldsymbol{\theta}_i, y_i) \quad \text{where} \quad \boldsymbol{\theta}_i = \tilde{\varphi}_w(x_i),$$

and $\ell(\boldsymbol{\theta}, y_i)$ is a loss function. Losses that are convex in $\boldsymbol{\theta}$ and lead to practically efficient algorithms have been proposed when (**CO-layer**) is linear on $\boldsymbol{\theta}$, which is the case on most applications. Typical examples include the structured Hinge loss [Nowozin, 2010] or the Fenchel-Young losses [Berthet et al., 2020]. The SPO+ loss proves successful when the training set contains target $\boldsymbol{\theta}_i$ instead of target y_i [Elmachtoub and Grigas, 2021].

In this paper, we focus on the learning by experience setting, where the training set (x_1, \dots, x_n) contains instances but not their solutions.

3.1 Learning problem and regularized learning problem

Let x_1, \dots, x_n be our *training set* composed of n instances of (Pb). Without loss of generality, we suppose that $c(y; x) \geq 0$ for all instances x and feasible solution $y \in \mathcal{Y}(x)$. We also suppose to have a mapping $u : x \mapsto u(x) \geq 0$ that is a coarse estimation of the absolute value of an optimal solution of x . We define the loss function as the weighted cost of the easy problem solution as a solution of the hard problem.

$$\ell(\mathbf{w}, x) := \frac{1}{u(x)} \max \left\{ c(y; x) : y \in \arg \min_{\tilde{y} \in \mathcal{Y}(x)} g(\tilde{y}, f_{\mathbf{w}}(x)) \right\}. \quad (9)$$

The *learning problem* consists in minimizing the expected loss on the training set

$$\min_{\mathbf{w} \in \mathbf{W}} \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{w}, x_i). \quad (10)$$

The instances in the training set may be of different size, leading to solutions costs which different order of magnitudes. The weight $\frac{1}{u(x)}$ enables to avoid giving too much importance to large instances.

When the approximation is flexible and the training set is small, the solution of (10) may overfit the training set, and lead to poor performance on instances that are not in the training set. In that case, the usual technique to avoid overfitting is to regularize the problem. One way to achieve this is to make the prediction “robust” with respect to small perturbations: We want the solution returned to be good even if we use $\mathbf{w} + \mathbf{Z}$ instead of \mathbf{w} , where \mathbf{Z} is a small perturbation. Practically, we assume that \mathbf{Z} is a standard Gaussian, $\sigma > 0$ is a real number, and we define the *perturbed loss*

$$\ell^{\text{pert}}(\mathbf{w}, x) = \mathbb{E}_{\mathbf{Z}} \left[\frac{1}{u(x)} \max \left\{ c(y; x) : y \in \arg \min_{\tilde{y} \in \mathcal{Y}(x)} g(\tilde{y}, f_{\mathbf{w} + \sigma \mathbf{Z}}(x)) \right\} \right]. \quad (11)$$

This perturbation can be understood as a regularization of the easy problem [Berthet et al., 2020]. The *regularized learning problem* is then formulated as follows.

$$\min_{\mathbf{w} \in \mathbf{W}} \frac{1}{n} \sum_{i=1}^n \ell^{\text{pert}}(\mathbf{w}, x) \quad (12)$$

3.2 Algorithms to solve the learning problem

Proposition 1. *If $\mathbf{w} \mapsto f_{\mathbf{w}}(x)$ and $\boldsymbol{\theta} \mapsto g(y, \boldsymbol{\theta})$ are piecewise linear for all y in $\mathcal{Y}(x)$, then the objective of (10) is piecewise constant in x .*

Proof. Since the composition of two piecewise linear functions is piecewise linear, $\mathbf{w} \mapsto g(x; f_{\mathbf{w}}(x))$ is piecewise linear. Hence, there exists a partition of the space into a finite number of polyhedra such that the set $\left\{ c(x; x) : x \in \arg \min_{x \in \mathcal{X}(\mathcal{E})} g(x; \varphi_{\mathbf{w}}(x)) \right\}$ is constant on each polyhedron. The definition of $\ell(\mathbf{w}, x)$ then ensures that $\mathbf{w} \mapsto \ell(\mathbf{w}, x)$ is piecewise constant on the interior of each polyhedron of the partition, and lower semi-continuous, which gives the result. \square

Proposition 1 is bad news from an optimization point of view. We need a black-box optimization algorithm that uses a moderate amount of function evaluations, does not rely on “slope” (due to null gradient), and takes a global approach (due to non-convexity). We therefore suggest using either a heuristic that searches the state space such as the DIRECT algorithm [Jones et al., 1993], or a Bayesian optimization algorithm that builds a global approximation of the objective function and uses it to sample the areas in the space of \mathbf{w} that are promising according to the approximation. The numerical experiments evaluate the performance of these two kinds of algorithms.

Let us now consider the regularized learning problem (12). Since the convolution product of two functions is as smooth as the most smooth of the two functions, $\mathbf{w} \mapsto \ell^{\text{pert}}(\mathbf{w}, x)$ is C^∞ . It can therefore be minimized using a stochastic gradient descent [Dalle et al., 2022]. On our applications, and using a generalized linear model, we obtained better results by solving a sample average approximation of this perturbed learning problem using the heuristics mentioned above. This is not so surprising because in that case, the objective of the learning problem is composed of several plateaus with smooth transition inbetween, which is not much easier to solve in practice. Remark that stochastic gradient descent is the method of choice when using a large neural network.

3.3 Practical remarks for a generic implementation

Perturbation strength. Section 4 provides a closed formula to set the perturbation strength σ .

Skipping the bilevel optimization Using a bilevel optimization enables to define $\ell(\mathbf{w}, x)$ unambiguously even when the easy problem (CO-layer) admits several optimal solutions. Since the bilevel optimization is not easy to handle, we use in practice the loss

$$\tilde{\ell}(\mathbf{w}, x) = \frac{1}{u(x)} c(\mathcal{A} \circ \varphi_{\mathbf{w}}(x))$$

that takes the solution returned by the algorithm \mathcal{A} we use for (CO-layer). Its value may therefore depend on \mathcal{A} .

Post-processing On many applications, the post-processing h is time-consuming, and there exists an alternative post-processing \tilde{h} that is much faster, even if the resulting solution z may have a larger cost. A typical example is our $1/|r_j| \sum_j C_j$ application, where $\mathcal{Y}(x) = \mathcal{Z}(x)$, and the post-processing is only a local descent. The post-processing is therefore not mandatory, and we could use $\tilde{h} = \text{Identity}$. In that context, using \tilde{h} instead of h during the learning phase leads to a much faster learning algorithm, while not necessarily hurting the quality of the \mathbf{w} learned.

Sampling in the prediction pipeline. If we are ready to increase the execution time, the perturbation of \mathbf{w} by \mathbf{Z} can also be used to increase the quality of the solution returned

by our solution pipeline. We can draw several samples \mathbf{Z}_i of \mathbf{Z} , apply the solution pipeline with $\mathbf{w} + \sigma \mathbf{Z}_i$ instead of \mathbf{w} , and return the best solution found across the samples at the end. We provide numerical results with this perturbed algorithm on the $1|r_j|\sum_j C_j$ problem in Section 5.

4 Learning rate and approximation ratio

This section introduces theoretical guarantees on the average optimality gap of the solution returned by the learned algorithm when \mathbf{w} is chosen as in Section 3. Two conditions seem necessary to obtain such guarantees. First, it must be possible to approximate the hard problem by the easy one. That is, there must exist a $\tilde{\mathbf{w}}$ such that an optimal solution of $\varphi_{\tilde{\mathbf{w}}}(x)$ provides a good solution of x . And second, when such a $\tilde{\mathbf{w}}$ exists, our learning problem must be able to find it or another \mathbf{w}' that leads to a good approximation. Our proof strategy is therefore in two steps. First, we show that the solution of our learning problem converges toward the “best” \mathbf{w} when the number of instances in the solution set increases. And then we show that if there exists a $\tilde{\mathbf{w}}$ such that the expected optimality gap of the solution returned by our solution approach is bounded, then the expected optimality gap for the learned \mathbf{w} is also bounded. For statistical reasons discussed at the end of Section 4.2, we carry this analysis using the regularized learning problem (12).

4.1 Background on learning with perturbed bounded losses

Let ξ be a random variable on a space Ξ and \mathbf{W} a non-empty compact subset of \mathbb{R}^d , $\mathbf{W} \subseteq \mathbb{B}_\infty(M)$ where $\mathbb{B}_\infty(M)$ is the $\|\cdot\|_\infty$ ball of radius M on \mathbb{R} . Let $\ell : \Xi \times \mathbb{R}^d \rightarrow [0, 1]$ be a loss function, we define the perturbed loss as

$$\ell^{\text{pert}}(\bar{\xi}, \mathbf{w}) = \mathbb{E}[\ell(\bar{\xi}, \mathbf{w} + \sigma Z)] \quad \text{with } \sigma > 0. \quad (13)$$

We suppose that $\ell(\cdot, \mathbf{w})$ is integrable for all $\mathbf{w} \in \mathbf{W}$. We define the *expected risk* $L(\mathbf{w})$ and the *expected risk minimizer* \mathbf{w}^* as

$$\mathbf{w}^* \in \arg \min_{\mathbf{w} \in \mathbf{W}} L(\mathbf{w}) \quad \text{with } L(\mathbf{w}) = \mathbb{E}[\ell^{\text{pert}}(\xi, \mathbf{w})]. \quad (14)$$

Let ξ_1, \dots, ξ_n be n i.i.d. samples of ξ . We define the *empirical risk* $\hat{L}_n(\mathbf{w})$ and the empirical risk minimizer $\hat{\mathbf{w}}_n$ as

$$\hat{\mathbf{w}}_n \in \arg \min_{\mathbf{w} \in \mathbf{W}} \hat{L}_n(\mathbf{w}) \quad \text{with } \hat{L}_n(\mathbf{w}) = \frac{1}{n} \sum_{i=1}^n \ell^{\text{pert}}(\xi_i, \mathbf{w}). \quad (15)$$

Note that both $L_n(\mathbf{w})$ and $\hat{\mathbf{w}}_n$ are random due to the sampling of the training set ξ_1, \dots, ξ_n . The following result bounds the excess risk incurred when we use $L_n(\mathbf{w})$ instead of $L(\mathbf{w})$.

Theorem 2. Suppose that $\mathbf{W} \subseteq \mathbb{B}_\infty(M)$ where $\mathbb{B}_\infty(M)$ is the $\|\cdot\|_\infty$ ball of radius M on \mathbb{R} . Given $0 < \delta < 1$, with probability at least $1 - \delta$, we have the following bound on the excess risk.

$$L(\hat{\mathbf{w}}_n) - L(\mathbf{w}^*) \leq C \frac{Md}{\sigma\sqrt{n}} + \sqrt{\frac{2\log(2/\delta)}{n}} \quad (16)$$

with $C = 48 \int_0^1 \sqrt{-\log x} dx$.

We believe that Theorem 2 is in the statistical learning folklore, but since we did not find a proof, we provide one based on classical statistical learning results in Appendix A.

Learning rate of our structured approximation In order to apply Theorem 2 to the learning problem of Section 3, we must endow the set \mathcal{X} of instances with a distribution. Recall that the loss ℓ and the perturbed loss ℓ^{pert} have been defined in Equations (9) and (11). We assume that the instance x is a random variable with probability distribution μ on \mathcal{X} , and that both $\ell(\cdot, \mathbf{w})$ and $\ell^{\text{pert}}(\cdot, \mathbf{w})$ are integrable for all \mathbf{w} . With these definitions, using \mathcal{X} as Ξ , instances x as random variables ξ , and if we suppose that the training set x_1, \dots, x_n is composed of n i.i.d. samples of x , we have recast our regularized learning problem (12) as a special case of (15). We can therefore apply Theorem 2 and deduce that the upper bound (16) on the excess risk applies.

This result underlines a strength of the architectures of Section 2.3. Because they enable to use approximations parametrized by \mathbf{w} whose dimension is small and does not depend on x , the bound on the excess risk only depends on the dimension of \mathbf{w} and not on the size of the instances used. Hence, *a pipeline with these architectures enables to make predictions that generalize (in expectation) on a test set whose instances structures $\mathcal{I}(x)$ are not necessarily present in the training set.* This is confirmed experimentally in Section 5, where the structures of the instances in the test set of the stochastic VSP (the graph D) do not appear in the training set.

Remark 1. Theorem 2 does not take into account the fact that, practically and in our numerical experiments, we use a sample average approximation on \mathbf{Z} of the perturbed loss instead of the true perturbed loss. \triangle

4.2 Approximation ratio of our structured approximation

Let $d(x) = |I(x)|$ and $c^*(x) = \arg \min_{y \in \mathcal{Y}(x)} c(y, x)$ be the cost of an optimal solution of (Pb).

Theorem 3. Suppose that for all x in \mathcal{X} (outside a negligible set for the measure on \mathcal{X}),

1. $f_{\mathbf{w}}(x) = (\langle \mathbf{w} | \boldsymbol{\phi}(i, x) \rangle)_{i \in \mathcal{I}(x)}$ and $\|\boldsymbol{\phi}(i, x)\|_2 \leq \kappa_\phi$,

2. and there exists $\tilde{\mathbf{w}}$, $a > 0$, $b > 0$, and $\beta \in \{1, 2\}$ such that, for any $\mathbf{p} \in \mathbb{R}^{d(x)}$,

$$c(y, x) - c^*(x) \leq au(x) + b\|\mathbf{p}\|_\beta \quad \text{for any } y \in \arg \min_{\tilde{y} \in \mathcal{Y}(x)} g(\tilde{y}, f_{\tilde{\mathbf{w}}}(x) + \mathbf{p})$$

Then, under the hypotheses of [Theorem 2](#), with probability at least $1 - \delta$ (on the sampling of the training set)

$$\underbrace{L(\hat{\mathbf{w}}_n) - \mathbb{E} \left[\frac{c^*(x)}{u(x)} \right]}_{\text{Perturbed prediction optimality gap}} \leq \underbrace{C \frac{Md}{\sigma\sqrt{n}} + \sqrt{\frac{2 \log(2/\delta)}{n}}}_{\text{Training set error}} + \underbrace{a}_{\text{Approximation error}} + \underbrace{b\sigma\kappa_\phi\sqrt{d} \mathbb{E} \left[\frac{[d(x)]^{1/\beta}}{u(x)} \right]}_{\text{Perturbation error}}$$

Before proving the theorem, let us make some comments. First, we explain why the hypotheses are meaningful. The first hypothesis only assumes that the model is linear, and that, with probability 1 on the choice of x in \mathcal{X} the features are bounded. Such a hypothesis is reasonable as soon as we restrict ourselves to instances whose parameters are bounded. Let us recall that $u(x)$ is a coarse upper bound on $c^*(x)$. When $u(x) = c^*(x)$ and $\mathbf{p} = 0$, the second hypothesis only means that our non-perturbed pipeline with parameter $\tilde{\mathbf{w}}$ is an approximation algorithm with ratio $1 + a$. With a $\mathbf{p} \neq 0$, the second hypothesis is stronger: It also ensures that this approximation algorithm guarantee does not deteriorate too fast. Later in this section, we prove that this hypothesis is satisfied for our running example.

Approximation ratio guarantee Using $u(x) = c^*(x)$ makes clear the fact that [Theorem 3](#) provides an approximation ratio guarantee in expectation. Furthermore, it gives a natural way of setting the strength σ of the perturbation: The bound is minimized when we use

$$\sigma_n = \sqrt{\frac{CM\sqrt{d}}{\sqrt{nb}\kappa_\phi \mathbb{E} \left[\frac{[d(x)]^{1/\beta}}{u(x)} \right]}}. \quad (17)$$

Using this optimal perturbation and $\delta = \frac{1}{n}$, the upper bound on $L(\hat{\mathbf{w}}_n) - \mathbb{E} \left[\frac{c^*(x)}{u(x)} \right]$ is in

$$a + O\left(n^{-1/4}(1 + \log(n))\right) \xrightarrow{n \rightarrow +\infty} a.$$

In other words, *in the large training set regime the learned $\hat{\mathbf{w}}_n$ recovers the approximation ratio guarantee a of $\tilde{\mathbf{w}}$.*

Large instances [Theorem 3](#) always provides guarantees when $d(x)$ is bounded on \mathcal{X} . However, it may fail to give guarantees when $d(x)$ is unbounded on \mathcal{X} since $\mathbb{E} \left[\frac{[d(x)]^{1/\beta}}{u(x)} \right]$ may not be finite. Since $u(x)$ is a coarse upper bound on $c^*(x)$, the term $\mathbb{E} \left[\frac{[d(x)]^{1/\beta}}{u(x)} \right]$ remains finite when $d(x)$ is unbounded only if the cost of an optimal solution $c^*(x)$ grows at least as fast as the number of parameters of the instance $d(x)$ to the power $1/\beta$. From that point of view, the single machine scheduling problem $1|r_j| \sum C_j$ is ideal. Indeed, in that case $d(x)$ is the number of job, and $c^*(x) \sim [d(x)]^2$, hence $\left[\frac{[d(x)]^{1/\beta}}{u(x)} \right]$ becomes smaller and smaller when the size of x increases. On the two stage spanning tree problem, the situation is slightly less favorable. Indeed, $d(x)$ is equal to twice the number of edges. Since an optimal spanning tree contains

$|V|-1$ edges, we expect $c^*(x)$ to be of the order of magnitude of $d(x)$ on sparse graphs (graphs such that $|E| \sim |V|$, like grids for instance), and $\sqrt{d(x)}$ on dense graphs (graph such that $|E| \sim |V|^2$, like complete graphs). Later in this section, we prove that the second hypothesis is satisfied for the two stage spanning tree problem with $\beta = 1$. Hence, Theorem 3 gives an approximation ratio guarantee for sparse graphs. The situation is roughly the same for the stochastic vehicle scheduling problem. A typical example where $\mathbb{E} \frac{[d(x)]^{1/\beta}}{u(x)}$ may not be finite is the shortest path problem on a dense graph. On many applications such as finding an optimal journey on a public transport system, the number of arcs tends to remain bounded, say ≤ 10 , while the number of arcs in the graph $d(x)$ grows with the size of the instance.

Optimal resolution of the learning problem Theorem 3 applies for the optimal solution $\hat{\mathbf{w}}_n$ of the learning problem. We let it to future work to design an exact algorithm that guarantees that the \mathbf{w} returned is within an optimality gap γ with the optimal solution. We would then obtain a variant of Theorem 3 proving the approximation ratio result for the \mathbf{w} returned, with an additional term in γ in the upper bound taking into account the optimality gap.

Influence of the perturbation Since we have made very few assumptions on c , g , and $f_{\mathbf{w}}$, we do not have control on the size of the family of functions $\{\ell_{\mathbf{w}} : \mathbf{w}\}$. This family may be very large, and therefore able to fit any noise, which would lead to slower learning rate. Without additional assumptions, we therefore need to regularize the family. In particular, we need to smooth the piecewise constant loss (Proposition 1). As we have seen in this section, perturbing \mathbf{w} does the job, but comes at a double cost in Theorem 3: a perturbation error, and larger than hoped training set error in $O(d/\sqrt{n})$. The term in $O(d/\sqrt{n})$ is slightly disappointing because the proof techniques used in statistical learning theory typically lead to bounds in $O(\sqrt{d}/\sqrt{n})$. This is for instance the case for the metric entropy method used to prove Theorem 2 when the gradient of the loss is Lipschitz in \mathbf{w} . The Gaussian perturbation restores the Lipschitz property for the perturbed loss, but it comes at the price of an additional \sqrt{d} in the bound derived by the metric entropy method, as can be seen in the proof of Lemma 9 in Appendix A. Designing a learning approach that avoids the additional \sqrt{d} term is an interesting open question. An alternative would be to make more assumptions on c , g , and φ with the objective of making the perturbation optional in the proof.

Proof of Theorem 3 We have

$$L(\hat{\mathbf{w}}_n) - \mathbb{E} \left[\frac{c^*(x)}{u(x)} \right] = \underbrace{L(\hat{\mathbf{w}}_n) - L(\mathbf{w}^*)}_{\leq C \frac{Md}{\sigma\sqrt{n}} + \sqrt{\frac{2 \log(2/\delta)}{n}}} + \underbrace{L(\mathbf{w}^*) - L(\tilde{\mathbf{w}})}_{\leq 0} + L(\tilde{\mathbf{w}}) - \mathbb{E} \left[\frac{c^*(x)}{u(x)} \right]$$

We therefore need to upper bound $L(\tilde{\mathbf{w}}) - \mathbb{E}\left[\frac{c^*(x)}{u(x)}\right]$ by $a + b\sigma\kappa_\phi\sqrt{d}\mathbb{E}\left[\frac{[d(x)]^{1/\beta}}{u(x)}\right]$. We have

$$\begin{aligned} \|f_{\tilde{\mathbf{w}}+\sigma\mathbf{Z}}(x) - f_{\tilde{\mathbf{w}}}(x)\|_\beta &= \left\| \sigma(\langle \mathbf{Z} | \boldsymbol{\phi}(i, x) \rangle)_{i \in \mathcal{I}(x)} \right\|_\beta \\ &\leq \left\| \sigma(\|\mathbf{Z}\|_2 \|\boldsymbol{\phi}(i, x)\|_2)_{i \in \mathcal{I}(x)} \right\|_\beta \\ &\leq \left\| \sigma(\|\mathbf{Z}\|_2 \kappa_\phi)_{i \in \mathcal{I}(x)} \right\|_\beta = \sigma\kappa_\phi \|\mathbf{Z}\|_2 [d(x)]^{1/\beta} \end{aligned}$$

Let y be in $\arg \min_{y \in \mathcal{Y}(x)} g(y, f_{\tilde{\mathbf{w}}+\sigma\mathbf{Z}}(x))$. The second hypothesis of the theorem and the previous inequality give

$$c(y, x) - c^*(x) \leq au(x) + b\kappa_\phi \|\mathbf{Z}\|_2 [d(x)]^{1/\beta}.$$

Since \mathbf{Z} is a standard Gaussian, we get $\mathbb{E}\|\mathbf{Z}\| \leq \sqrt{d}$ (Equation (23) in Appendix A), and the result follows by dividing the previous equality by $u(x)$ and taking the expectation. \square

4.3 Existence of a $\tilde{\mathbf{w}}$ with an approximation ratio guarantee

In this section, we prove that the hypotheses of Theorem 3 are satisfied for the maximum weight two stage spanning tree problem. We then give a criterion which ensures that these hypotheses are satisfied.

Maximum weight two stage spanning tree In this section, we restrict ourselves to maximum weight spanning tree instances, that is, instances of the minimum weight spanning tree with $c_e \leq 0$ and $d_{es} \leq 0$ for all e in E and s in S . Given an instance, let $I(x) = \{(e, \text{stage}) : e \in E, \text{stage} \in \{1, 2\}\}$, leading to $\bar{c}_e = \langle \mathbf{w} | \boldsymbol{\phi}((e, 1), x) \rangle$ and $\bar{d}_e = \langle \mathbf{w} | \boldsymbol{\phi}((e, 2), x) \rangle$. We define a feature

$$\phi(e, 1) = c_e \quad \text{and} \quad \phi(e, 2) = \frac{1}{|S|} \sum_{s \in S} d_{es},$$

and define $\tilde{\mathbf{w}}$ to be equal to 1 for this feature and 0 otherwise. The following proposition shows that the second hypothesis of Theorem 3 is then satisfied with $a = 1/2$ and $b = 1$.

Proposition 4. *For any instance x of the maximum weight spanning tree problem, we have*

$$c(y, x) - c^*(x) \leq \frac{1}{2}|c^*(x)| + \|\mathbf{p}\|_1 \quad \text{for any } y \in \arg \min_{y' \in \mathcal{Y}(x)} g(y, f_{\tilde{\mathbf{w}}}(x) + \mathbf{p}).$$

Our proof shows that the results stands for $a = \frac{|S|-1}{2|S|-1}$ when $|S|$ is upper-bounded by M . Combined with Theorem 3, Proposition 4 ensures that, when the training set is large, the learned \hat{w}_n has the approximation ratio guarantee proved by Escoffier et al. [2010] for the two stage maximum weight spanning tree. The proof of Proposition 4 is an extension of the proof of Escoffier et al. [2010, Theorem 6] to deal with non-zero perturbations \mathbf{p} .

Proof of Proposition 4. The proof will use the following well known result

Lemma 5. *Let $x \mapsto f_1(y)$ and $y \mapsto f_2(y)$ be functions from compact set K to \mathbb{R} , and let y_1^* and y_2^* be respectively minima of f_1 and f_2 . If we have $|f_1(y) - f_2(y)| \leq \gamma$ for all y , then $f_1(y_2^*) - f_1(y_1^*) \leq 2\gamma$.*

We fix an instance x . Let us first introduce some solutions of interest. Given a θ , let us denote by $\bar{y}(\theta) = (\bar{E}_1(\theta), \bar{E}_2(\theta))$ the result of the prediction problem (2). Let $\bar{z}(\theta) = (\bar{E}_1(\theta), (\bar{E}_s(\theta)))$ with $\bar{E}_s(\theta) = \bar{E}_2(\theta)$. Let $\hat{z}(\theta) = (\bar{E}_1(\theta), (\hat{E}_s(\theta)))$ where $E_s(\theta)$ is the optimal second stage decision for scenario s when the first stage decision is $E_1(\theta)$

$$\hat{E}_s(\theta) \in \arg \min \left\{ \sum_{e \in E_s} d_{es} : E_s \subseteq E, E_s \cap \bar{E}_1(\theta) = \emptyset, (V, E_s \cup \bar{E}_1(\theta)) \in \mathcal{T} \right\}. \quad (18)$$

We denote by $z^\emptyset = (\emptyset, (E_s^\emptyset))$ the solution with no first stage: (V, E_s^\emptyset) is a minimum weight spanning tree for second stage weights $(d_{es})_s$ of scenario s . And finally, we denote by $z(\theta) = (E_1, (E_s)_{s \in S})$ the solution returned by our pipeline, which is the solution of minimum cost among $\hat{z}(\theta)$ and z^\emptyset .

Let $z^* = (E_1^*, (E_s^*))$ be an optimal solution of (1), and $\bar{y}^{s,*} = (E_1^*, E_s^*)$ be the solution of (2) obtained by taking E_1^* as first stage solution and E_s^* as second stage solution.

First, consider solution $\bar{z} = (\bar{E}_s, (\bar{E}_s)_s)$ of (1) such that the second stage solution \bar{E}_s is identical and equal to \bar{E}_2 for all scenarios s in S , and denote by $\bar{y} = (\bar{E}_1, \bar{E}_2)$ the solution of (2) obtained by taking \bar{E}_1 as first stage solution and \bar{E}_2 as second stage solution. It follows from the definition of \tilde{w} that

$$C(\bar{z}) = g(\bar{y}, \tilde{\theta}) \quad \text{where} \quad \tilde{\theta} = f_{\tilde{w}}(x). \quad (19)$$

Second, remark that, for any θ , $y = (E_1, E_2)$, and \mathbf{p} in $\mathbb{R}^d(x)$, we have

$$|g(y, \theta) - g(y, \theta + Z)| = |\langle Z | y \rangle| = \sum_{e \in E_1} p_{e1} + \sum_{e \in E_2} p_{e2} \leq \|\mathbf{p}\|_1 \quad (20)$$

where the last inequality comes from the fact that y is the indicator vector of a tree.

Let s be scenario in S . We have

$$\begin{aligned} C(\hat{z}(\tilde{\theta} + Z)) &\leq C(\bar{z}(\tilde{\theta} + Z)) && \text{Optimal second stage} \\ &= g(\bar{y}(\tilde{\theta} + Z), \tilde{\theta}) && \text{Equation (19)} \\ &\leq g(\bar{y}(\tilde{\theta}), \tilde{\theta}) + 2\|\mathbf{p}\|_1 && \text{Equation (20) + Lemma 5} \\ &\leq g(\bar{y}^{s,*}, \tilde{\theta}) + 2\|\mathbf{p}\|_1 && \text{Optimality of } \bar{y}(\tilde{\theta}) \\ &= \sum_{e \in \bar{E}_1^*} c_e + \frac{1}{|S|} \sum_{e \in \bar{E}_s^*} \sum_{s' \in S} d_{es'} + 2\|\mathbf{p}\|_1 \\ &\leq \sum_{e \in \bar{E}_1^*} c_e + \frac{1}{|S|} \sum_{e \in \bar{E}_s^*} d_{es} + 2\|\mathbf{p}\|_1 && d'_{es} \leq 0 \text{ for all } e, s'. \end{aligned}$$

Furthermore, since (V, E_s^\emptyset) is a minimum spanning tree with $(d_{es})_e$ edge weights,

$$C(z^\emptyset) = \frac{1}{|S|} \sum_{s \in S} \sum_{e \in E_s^\emptyset} d_{es} \leq \frac{1}{|S|} \sum_{s \in S} \left[\sum_{e \in E_1^*} \underbrace{d_{es}}_{\leq 0} + \sum_{e \in E_s^*} d_{es} \right] \leq \frac{1}{|S|} \sum_{s \in S} \sum_{e \in E_s^*} d_{es}.$$

Summing the two previous inequalities, we get

$$\begin{aligned} C(z(\boldsymbol{\theta})) &= \min(C(\hat{z}(\tilde{\boldsymbol{\theta}} + Z)), C(z^\emptyset)) \\ &\leq \frac{|S|C(\hat{z}(\tilde{\boldsymbol{\theta}} + Z)) + (|S| - 1)C(z^\emptyset)}{2|S| - 1} \\ &\leq \frac{|S|(\sum_{e \in E_1^*} c_e + 2\|\mathbf{p}\|_1) + \sum_s \sum_{e \in E_s^*} d_{es}}{2|S| - 1} \\ &= \frac{|S|}{2|S| - 1} (C(z^*) + 2\|\mathbf{p}\|_1) \leq \frac{1}{2}C(z^*) + \|\mathbf{p}\|_1 \end{aligned}$$

Since $C(z^*) \leq 0$, we get $C(z(\boldsymbol{\theta})) - C(z^*) \leq \frac{1}{2}|C(z^*)| + \|\mathbf{p}\|_1$, which is the result searched. \square

Remark that we have proved the stronger bound $c(y, x) - c^*(x) \leq \frac{|M|-1}{2|M|-1}|c^*(x)| + \|\mathbf{p}\|_1$ when $|S|$ is upper bounded by M on \mathcal{X} .

Objective function approximation Let us finally remark that the hypotheses of Theorem 3 are satisfied when $g(y, f_{\mathbf{w}}(x))$ is a good approximation of $\boldsymbol{\theta} \mapsto g(y, \boldsymbol{\theta})$ for some \mathbf{w} .

Lemma 6. *Suppose that $\boldsymbol{\theta} \mapsto g(y, \boldsymbol{\theta})$ is κ_g Lipschitz in $\|\cdot\|_\beta$, and there exists $\tilde{\mathbf{w}} \in \mathbf{W}$ and $\alpha > 0$ is such that, for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}(x)$, we have*

$$\frac{|c(y, x) - g(y, f_{\tilde{\mathbf{w}}}(x))|}{u(x)} \leq \alpha. \quad (21)$$

Then the second hypothesis of Theorem 3 is satisfied with $a = 2\alpha$ and $b = 2\kappa_g$.

Proof. For any y in $\mathcal{Y}(x)$, we have

$$\begin{aligned} &|c(y, x) - g(x, f_{\tilde{\mathbf{w}}}(x) + \mathbf{p})| \\ &\leq |c(y, x) - g(x, f_{\tilde{\mathbf{w}}}(x))| + |g(x, f_{\tilde{\mathbf{w}}}(x)) - g(x, f_{\tilde{\mathbf{w}}}(x) + \mathbf{p})| \\ &\leq \alpha u(x) + \kappa_g \|\mathbf{p}\|_\beta \end{aligned}$$

The results therefore follows from Lemma 5 and the previous inequality. \square

5 Numerical experiments

This section tests the performance of our algorithms on our running example, the stochastic vehicle scheduling problem, and the $1|r_j| \sum_j C_j$ scheduling problem of Section 2. For the

two latter applications, we use the same encoding f_w , easy problem solution algorithm, and decoding ψ as in previous contributions [Parmentier, 2021, Parmentier and T’Kindt, 2021]. The only difference is that, instead of using the learning by demonstration approaches proposed in these papers, we use the learning by experience approach of this paper. All the numerical experiments have been performed on a Linux computer running Ubuntu 20.04 with an Intel® Core™ i9-9880H CPU @ 2.30GHz \times 16 processor and 64 GiB of memory. All the learning problem algorithms are parallelized: The value of the loss on the different instances in the training set are computed in parallel. The prediction problem algorithms are not parallelized.

5.1 Maximum weight two stage spanning tree

Let us now consider the performance of our pipeline on the maximum weight two stage spanning tree. All the algorithms are implemented in `julia`. The code to reproduce the numerical experiments is open source¹.

Training, validation and test sets We use instances on square grid graphs of width $\{10, 20, 30, 40, 50, 60\}$, i.e., with $|V|$ in $\{100, 400, 900, 1600, 2500, 3600\}$. First stage weights are uniformly sampled on the integers in $\{-20, \dots, 0\}$. Second stage weights are uniformly sampled on the integers in $\{-K, \dots, 0\}$ with $K \in \{10, 15, 20, 25, 30\}$. Finally, instances have 5, 10, 15 or 20 second stage scenarios. Our training set, validation set, and test set contain 5 instances for each grid width, weight parameter K , and number of scenarios. The training, validation and test sets therefore each contain 600 instances.

Bounds and benchmarks On each instance of the training, validation, and test set, we solve the Lagrangian relaxation problem using a subgradient descent algorithm for 50,000 iterations, which provides a lower bound on an optimal solution. We also run a Lagrangian heuristic based on the final value of the duals.


We use three benchmarks to evaluate our algorithms: the Lagrangian heuristic, the approximation algorithm of Escoffier et al. [2010], and our pipeline trained by imitation learning using a Fenchel Young loss to reproduce the solution of the Lagrangian heuristic. Remark that since the approximation algorithm, the pipeline learning by imitation, and the pipeline learned with our loss are all instances of our pipeline, they take roughly the same time. On the contrary the Lagrangian heuristic requires to solve the 50,000 iterations of the subgradient descent algorithm, and is therefore 4 order of magnitude slower.

Hyperparameters tuning We use a sample average approximation of our perturbed loss ℓ^{pert} with 20 scenarios. Figure 2.a provides the average value of the gap between the solution returned by our pipeline and the Lagrangian lower bound on the validation set for the model learned with different value of ε . Based on these results, we use $\varepsilon = 0.001$.

¹<https://github.com/axelparmentier/MaximumWeightTwoStageSpanningTree.jl>

ε	Gap
0.0e+00	2.8%
1.0e-04	2.7%
3.0e-04	2.7%
1.0e-03	2.7%
3.0e-03	2.7%
1.0e-02	2.7%
3.0e-02	3.9%
1.0e-01	5.5%
3.0e-01	59.5%

(a)



(b)

Figure 2: Maximum weight spanning tree. (a) Hyperparameters tuning on the validation dataset for the model learned on the training set. b, Gap with respect to the Lagrangian relaxation bound as a function of (b) the number of vertices $|V|$ with $\varepsilon = 0.001$.

Results Figure 2.b illustrates the average and worst gap with respect to the Lagrangian bound obtained on the test set for our pipeline and the different benchmarks. The pipeline learning by imitation of by experience enable to match the performance of the Lagrangian heuristic, while being 4 order of magnitude faster. These three algorithms significantly outperform the approximation algorithm. In summary, our pipeline learned by experience enables to retrieve the performance of the best algorithms as well as the theoretical guarantee of the approximation algorithm.

5.2 Stochastic VSP

5.2.1 Setting: Features, post-processing, and instances

For the numerical experiments on the stochastic VSP, we use the exact same settings as in our previous work [Parmentier, 2021]. We use the same linear predictor with a vector ϕ containing 23 features. And we do not use a post-processing ψ . The easy problem is solved with Gurobi 9.0.3 using the LP formulation based on flows.

We also use the same instance generator. This generator takes in input the number of tasks $|V|$, the number of scenarios $|\Omega|$ in the sample average approximation, and the seed of the random number generator. We say that an instance is of moderate size if $|V| \leq 100$, of large size if $100 \leq |V| \leq 750$, and of huge size if $1000 \leq |V|$. Table 2 summarizes the instances generated. The first two columns indicate the size of the instances. A \checkmark in the next five columns $|\Omega|$ indicates that instances with $|\Omega|$ scenarios are generated for instance size $|V|$ considered. The last five columns detail the composition of the different sets of instances: Three training sets, one validation set (Val), and a test set (Test). The table can be read as follows: The training set (small) contains $10 \times |\{50, 100, 200, 500, 1000\}| = 50$ instances, each of these having 50 tasks in V , but no larger instances. The test set contains

Size	$ V $	$ \Omega $					Train (small)	Train (moderate)	Train (all)	Val	Test
		50	100	200	500	1000					
Moderate	50	✓	✓	✓	✓	✓	10	5	1	2	8
	75	✓	✓	✓	✓	✓		5	1	2	8
	100	✓	✓	✓	✓	✓		5	1	2	8
Large	200	✓	✓	✓	✓	✓	5	1	2	8	
	500	✓	✓	✓	✓	✓		1	2	8	
	750	✓	✓	✓	✓	✓		1	2	8	
Huge	1000	✓	✓	✓	✓	✓	1	2	8		
	2000	✓						2	8		
	5000	✓						2	8		

Table 2: Instances considered for the stochastic VSP.

instances of all size. For instance, it contains $8 \times |\{50, 100, 200, 500, 1000\}| = 40$ instances of size 50 and 8 instances of size 5000. For the largest sizes, we use only instances with 50 scenarios for memory reasons: The instances files already weigh several gigabytes.

The small training set, the validation set, and the test set are identical to those previously used [Parmentier, 2021]. The validation set, which is used in the learning by demonstration approach, is not used on the learning by experience approach, since we do not optimize on classifiers hyperparameters. This previous contribution considers only the “small” training set, with 50 instances with 50 tasks, it uses a learning by demonstration approach and exact solvers cannot handle larger instances. This is no more a constraint with the learning by experience approach proposed in this paper. We therefore introduce two additional training sets: one that contains 100 instances of moderate size, and one containing 35 instances of all sizes. These training sets are relatively small in terms of number of instances, but they already lead to significant learning problem computing time and good performance on the test set.

5.2.2 Learning algorithm

On each of the three training sets, we solve the learning problem (10) and the regularized learning problem (12). We use the number of tasks $|V|$ as $u(x)$. It is not an upper bound on the cost, but the cost of the optimal solution scales almost linearly with $|V|$. In both case, we solve the learning problem on the L_∞ ball of radius 10. For the regularized learning problem, we use a perturbation strength of intensity $\sigma = 1$, and we solve the sample averaged approximation of the problem with 100 scenarios. We evaluate two heuristic algorithms: The DIRECT algorithm [Jones et al., 1993] implemented in the `nlopt` library [Johnson], and the Bayesian optimization algorithm as it is implemented in the `bayesopt` library [Martinez-Cantin]. We run each algorithm on 1000 iterations, which means that they can compute the objective function 1000 times. Both algorithms are launched with the default parameters of the libraries. In particular, the Bayesian optimization algorithm uses the anisotropic kernel

Learning problem			DIRECT		Bayes Opt	
Obj.	Train. set	pert	CPU time (hh:mm:ss)	Obj	CPU time (days, hh:mm:ss)	Obj
ℓ	small	–	0:01:20	290.39	0:09:06	287.48
	moderate	–	0:10:56	256.39	0:18:35	259.88
	all	–	2:56:40	231.66	3:56:18	235.04
ℓ^{pert}	small	100	0:11:52	286.11	0:37:22	287.95
	moderate	100	1:58:50	258.76	2:09:17	259.13
	all	100	22:44:03	233.84	1 day, 5:35:30	235.10

Table 3: Performance of the DIRECT [Johnson] and Bayesian optimization [Martinez-Cantin] on the learning problems (10) and (12) for the stochastic VSP.

with automatic relevance determination `kSum(kSEARD, kConst)` of the library.

Table 3 summarizes the result obtained with both algorithms. The first column contains the loss used: ℓ for the non-regularized problem (10) and ℓ^{pert} for the regularized problem. The next one provides the training set used. And the third column provides the number of samples used in the sample average approximation of the perturbation. The next four columns give the total computing time for the 1000 iterations and the value of the objective of the learning problem obtained at the end using the DIRECT algorithm and the Bayesian optimization algorithm.

The DIRECT algorithm approximates the value of the function based on a division of the space into hypercubes. At each iteration, the function is queried in the most promising hypercube, and the result is used to split the hypercube. The algorithm leverages a tractable lower bound to identify the most-promising hypercube and the split with few computations. Hence, the algorithm is very fast if the function minimized is not computationally intensive. The Bayesian optimization algorithm builds an approximation of the function minimized: It seeks the best approximation of the function in a reproducing kernel Hilbert space (RKHS) given the data available. At each iteration, it minimizes an activation function to identify the most promising point according to the model, evaluate the function at that point, and updates the approximation based on the value returned. Each of these steps are relatively intensive computationally. Hence, if the function minimized is not computationally intensive, the algorithm will be much slower than the DIRECT algorithm. This is what we observe on the first line of Table 3. Furthermore, in `bayesopt`, the DIRECT algorithm of `nlopt` is used to minimize the activation function. Our numerical experiments tend to indicate that the approximation in a RKHS does not enable to find a better solution than the simple exploration with DIRECT after 1000 iterations. Using Bayesian optimization may however be useful with a smaller iteration budget. Figure 5.2.2 provides the evolution of the objective function along time for the Bayesian optimization algorithm and the DIRECT algorithm on the learning problem corresponding to the last line of Table 3.

Since the DIRECT algorithm gives the best performance on most cases, we keep the \mathbf{w} returned by this algorithm for the numerical experiments on the test set. The performance




Figure 3: Learning problem objective value evolution as a function of time on stochastic VSP learning problem with perturbed loss ℓ^{pert} on training set “all”.

Learning problem	w	Obj	Train. set	Pert	Moderate			Large			Huge			All				
					T^{avg}	$\frac{T^{\text{max}}}{T^{\text{avg}}}$	δ^{avg}	δ^{max}	T^{avg}	$\frac{T^{\text{max}}}{T^{\text{avg}}}$	δ^{avg}	δ^{max}	T^{avg}	$\frac{T^{\text{max}}}{T^{\text{avg}}}$	δ^{avg}	δ^{max}	T^{avg}	$\frac{T^{\text{max}}}{T^{\text{avg}}}$
CRF	small	–	0.03	2.14	9.47%	20.47%	1.21	1.47	2.58%	6.77%	27.69	1.09	1.27%	1.88%	5.74	2.14	5.13%	20.47%
FYL	small	100	0.03	2.02	1.67%	4.23%	0.97	1.52	0.70%	2.10%	19.20	1.15	0.26%	1.06%	4.04	2.02	1.01%	4.23%
ℓ	small	–	0.03	2.33	4.37%	10.35%	0.82	1.47	3.65%	5.56%	20.42	1.26	3.29%	4.90%	4.21	2.33	3.88%	10.35%
	moderate	–	0.03	1.77	0.31%	3.24%	0.86	1.39	1.09%	2.92%	17.48	1.13	2.85%	6.18%	3.67	1.77	1.10%	6.18%
	all	–	0.03	1.56	0.52%	1.98%	0.84	1.33	0.07%	0.86%	18.10	1.14	0.07%	0.66%	3.78	1.56	0.25%	1.98%
ℓ^{pert}	small	100	0.03	1.81	2.90%	6.71%	0.84	1.53	2.55%	4.48%	16.29	1.99	2.04%	3.61%	3.44	1.99	2.59%	6.71%
	moderate	100	0.03	1.98	1.56%	5.00%	0.86	1.42	0.77%	2.12%	16.79	1.09	0.90%	1.82%	3.54	1.98	1.11%	5.00%
	all	100	0.03	1.63	1.05%	3.57%	0.87	1.45	1.10%	2.83%	18.08	1.30	1.16%	2.22%	3.79	1.63	1.09%	3.57%

The best results are in bold. CRF = Conditional Random Field

Table 4: Performance of our solution algorithm with different w on the stochastic VSP test set.

of the DIRECT algorithm could be improved using the optimization on the seed that will be introduced in Section 5.3.2.

5.2.3 Algorithm performance on test set

We now evaluate the performance of our solution pipeline with the w learned. It has been shown [Parmentier, 2021] that, using solution pipeline with the w learned by the structured learning approach with a conditional random field (CRF) loss on the small training set gives a state-of-the-art algorithm for the problem (the paper uses the maximum likelihood terminology instead of CRF loss). We therefore use it as a benchmark of the problem. We have also introduced a new learning by demonstration approach on the problem: We implement the Fenchel Young loss (FYL) structured learning approach [Berthet et al., 2020, Parmentier and T’Kindt, 2021] to obtain a second benchmark.

Table 4 summarizes the results obtained. The first three columns indicate how w has been computed: They provide the loss minimized as objective of the learning problem (Obj), the training set used, and for the approaches that use a perturbation, the number of scenarios

used in the sample average approximation (SAA).

The next columns provide the results on the test set. These columns are divided into four blocks giving results on the subsets moderate, large, huge instances of the test set and on the full test set. On each of these subsets of instances, we provide four statistics. The statistic T^{avg} provides the average computing time for our full solution pipeline on the subset of instances considered, which includes the computation of the features and $\varphi_{\mathbf{w}}(x)$, and the resolution of the easy problem with the LP solver (no decoding ψ is used). Most of this time is spent in the LP solver. Then, for each instance in the training set, we compute the ratio of the computing time for the instance divided by the average computing time for all the instances of the test set with the same number of tasks $|V|$. Indeed, we expect instances with the same $|V|$ to be of comparable difficulty. The column $\frac{T^{\text{max}}}{T^{\text{avg}}}$ gives the maximum value of this ratio on the subset of instances considered. Since we do not have an exact algorithm for the problem, for each instance x we compute the gap

$$\frac{c_{\mathbf{w}} - c^{\text{best}}}{c^{\text{best}}} \tag{22}$$

between the cost $c_{\mathbf{w}}$ of the solution returned by our solution pipeline with the \mathbf{w} evaluated and the cost of the best solution found for these instances using all the algorithms tested. The columns δ^{avg} and δ^{max} respectively provide the average and the maximum value of this gap on the set of instances considered. The two first lines provide the result obtained with the learning by demonstration benchmarks, and the next six ones obtained with the \mathbf{w} obtained with the learning algorithms of Table 3.

We can conclude from these experiments that:

1. When using the learning by demonstration approach, the Fenchel Young loss leads to better performances than the conditional random field loss.
2. Our learning by experience formulation gives slightly weaker performances than the learning by demonstration approach with a Fenchel Young loss when using the same training set.
3. Our learning by experience approach enables to use a more diversified training set, which enables it to outperform all the previously known approaches. The more diversified the training set, the better the performance.
4. The regularization by perturbation used tends to decrease the performance of the algorithm. This statement may no longer hold if we optimized the strength of the perturbation using a validation set.

5.3 Single machine scheduling problem $1|r_j|\sum_j C_j$

5.3.1 Setting

We use the exact same setting as the previous contribution on this problem [Parmentier and T'Kindt, 2021]. In particular, that paper introduces a vector of 66 features, and a

	Subsets of instances		
	Moderate	Large	Huge
Size n of instances in subset	{50, 75, 100, 150}	{200, 300, 500, 750}	{1000, 1500, 2000, 3000}

Table 5: Size of the $1|r_j|\sum_j C_j$ instances in the subsets of the test set.

subset of 27 features that leads to better performances. With the objective of testing what our learning algorithm can do on a larger dimensional problem, we focus ourselves on the problem with 66 features. And we also use the four kinds of decoding algorithms in that paper: no decoding (no ψ), a local search (LS), the same local search followed by release date improvement (RDI) algorithm (RDI \circ LS), and the perturbed versions of the last algorithm, (pert RDI \circ LS) where the solution pipeline is applied with $\mathbf{w} + \mathbf{Z}$ for 150 different samples of a standard Gaussian \mathbf{Z} , and keep the best solution found. RDI is a classic heuristic for scheduling problems, which is more time consuming but more efficient than the local search.

We use the same generator of instances as previous contributions [Della Croce and T'kindt, 2002, Parmentier and T'Kindt, 2021]. For a given instance with n jobs, processing times p_j are drawn at random following the uniform distribution $[1; 100]$ and release dates r_j are drawn at random following the uniform distribution $[1; 50.5n\rho]$. Parameter ρ enables to generate instances of different difficulties: We consider $\rho \in \{0.2, 0.4, 0.6, 0.8, 1.0, 1.25, 1.5, 1.75, 2.0, 3.0\}$. For each value of n and ρ , N instances are randomly generated leading for a fixed value of n to $10N$ instances. For the learning by demonstration approach, we use the same training set as [Parmentier and T'Kindt, 2021] with $n \in \{50, 70, 90, 110\}$ and $N = 100$, leading to a total of 4000 instances. We do not use larger instances because we do not have access to optimal solutions for larger instances. For the learning by experience approach, we use $n \in \mathcal{N} := \{50, 75, 100, 150, 200, 300, 500, 750, 1000, 1500, 2000, 3000\}$ and $N = 20$, leading to 2400 instances. In the test set, we use a distinct set of 2400 instances with $n \in \mathcal{N}$ and $N = 20$. This test set is almost identical to the one used in the literature [Parmentier and T'Kindt, 2021], the only difference being that instances with $n = 2500$ have been replaced by instance with $n = 3000$ to get a more balanced test set. Table 5 shows how we have partitioned this test set by number of jobs n in the instances, to get sets of instances of moderate, large, and huge size.

5.3.2 Learning algorithm

We use $n(n+1)$ as $u(x)$. It is not an upper bound on the cost, but the cost of the optimal solution scales roughly linearly with $u(x)$. We draw lessons from the stochastic VSP and use only a diverse training set of 4000 instances of all size in the training set. And because our solution pipeline for $1|r_j|\sum_j C_j$ is much faster than the one for the stochastic vehicle scheduling problem, we can use a larger training set. And we introduce two new perspectives. First, the DIRECT algorithm uses a random number generator. We observed that its performance is very dependent on the seed of the random number generator, and that using

obj	iter	pert	ψ	Tot. CPU	Avg \hat{L}	Best \hat{L}
ℓ	1000	–	–	0:05:46	36.71	35.19
ℓ	2500	–	–	0:12:31	36.64	35.29
ℓ^{pert}	1000	100	–	9:09:06	36.76	35.27
ℓ^{pert}	2500	100	–	20:45:52	36.69	35.25
ℓ	1000	–	LS	0:52:54	35.13	35.06
ℓ	2500	–	LS	1:47:42	35.12	35.06
ℓ^{pert}	1000	100	LS	3 days, 15:14:42	35.12	35.06
ℓ^{pert}	2500	100	LS	7 days, 15:24:08	35.11	35.06

Tot. CPU is given in days, hh:mm:ss.

Table 6: Learning algorithm results on $1|r_j|\sum_j C_j$.

a larger number of iterations does not necessarily compensate for the poor performance that would come from a bad seed. We therefore launch the algorithm with 10 different seeds, each time with a 1000 iterations budget, and report the best result. Second, as underlined in Section 3.3, it can be natural to use the loss ℓ^ψ where, instead of using the output of the easy problem, we use the output of the post-processing ψ . In our case, the post-processing is in two steps: first the local search, second the RDI heuristic. Since RDI is time-consuming, using it would lead to very large computing times on the training set used. We therefore take the solution at the end of the local search.

Table 6 summarizes the results obtained. The first column indicate if the perturbed loss or the non-perturbed loss has been used. The second indicates the number of iterations of DIRECT used. The third column indicates the number of scenarios used in the sample average approximation when the perturbed loss is used. And “–” (resp. LS) in the fourth column indicates if no (resp the local search) post-processing has been applied to the solution used in the loss. The column Tot. CPU then provides the total CPU time of the 10 runs of DIRECT with different seeds. Finally, the columns Avg \hat{L} and Best \hat{L} give respectively the average and the best loss value of the best solution found by DIRECT algorithm on the 10 seeds used.

We can conclude from these results that optimizing on the seed seems a good idea. We also observe that the loss function after the local search is smaller, which is natural given that the local search improves the solution found by the easy problem.

5.3.3 Algorithm performance on test set

Table 7 summarizes the results obtained with the different \mathbf{w} on the full test set. The first line corresponds to the Fenchel young loss (FYL) of the learning by demonstration approach previously proposed [Parmentier and T’Kindt, 2021], and serves as a benchmark. The next eight ones correspond to the parameters obtained solving the learning problem described in this paper with the settings of Table 6. The first four columns describe the parameters of the learning problems used to obtained \mathbf{w} and are identical to those of Table 6. The next columns indicate the average results on the full test set for the four kind of post-processing

Learning problem \mathbf{w}				Test set results (with several ψ)							
obj	iter	pert	ψ	no ψ		LS		RDI \circ LS		pert RDI \circ LS	
				δ^{avg}	δ^{max}	δ^{avg}	δ^{max}	δ^{avg}	δ^{max}	δ^{avg}	δ^{max}
FYL	–	–	–	1.81%	8.57%	1.10%	6.88%	0.07%	3.41%	0.02%	0.46%
ℓ	1000	–	–	0.63%	24.53%	0.34%	4.59%	0.06%	1.65%	0.02%	1.53%
ℓ	2500	–	–	1.08%	21.19%	0.30%	6.33%	0.07%	1.71%	0.04%	1.71%
ℓ^{pert}	1000	100	–	0.83%	23.61%	0.38%	3.64%	0.06%	1.65%	0.03%	1.37%
ℓ^{pert}	2500	100	–	0.75%	19.54%	0.33%	3.98%	0.06%	1.69%	0.02%	1.37%
ℓ	1000	–	LS	10.51%	54.67%	0.02%	1.30%	0.01%	1.12%	0.01%	1.12%
ℓ	2500	–	LS	10.16%	55.70%	0.02%	1.30%	0.01%	1.12%	0.01%	1.12%
ℓ^{pert}	1000	100	LS	10.54%	55.22%	0.03%	2.26%	0.02%	2.26%	0.02%	2.26%
ℓ^{pert}	2500	100	LS	10.51%	53.64%	0.03%	2.26%	0.02%	2.26%	0.02%	2.26%

Table 7: Performance of our solution algorithms with different \mathbf{w} on the $1|r_j|\sum_j C_j$ test set.

described in Section 5.3.1. Again, we provide the average δ^{avg} and the worse δ^{max} values of the gap (22) between the solution found by the algorithm and the best solution found by all the algorithms.

Two conclusions can be drawn from these results:

1. The solution obtained with our loss by experience approach tend to outperform on average those obtained using the Fenchel Young loss, but tend to have a poorer worst case behavior.
2. Using the loss with post-processing tend to improve the performance on the test set with the pipelines that use this preprocessing, and possibly other after. But it decreases the performance on the pipeline which do not use it.

Finally, Table 8 details the results for the fastest (no ψ) and the most accurate one (pert RDI \circ LS) pipeline on the subsets of instances of moderate, large, and huge size. In addition to the gaps, the average computing time T^{avg} is provided. Again, we can observe that:

3. Because the learning by experience approach enables to use a diversified set of instances in the training set, it outperforms the learning by demonstration approach on large and huge instances.

6 Conclusion

We have focused on heuristic algorithms for hard combinatorial optimization problems based on machine learning pipelines with a simpler combinatorial optimization problem as layer. Previous contributions in the literature required training sets with instances and their optimal solutions to train such pipelines. We have shown that the solutions are not necessarily needed, and we can learn such pipelines by experience if we formulate the learning problem

Pred. ψ	\mathbf{w}			Moderate			Large			Huge			
	obj	iter	$ \Omega $	ψ	T^{avg}	δ^{avg}	δ^{max}	T^{avg}	δ^{avg}	δ^{max}	T^{avg}	δ^{avg}	δ^{max}
no ψ	FYL	–	–	–	0.01	1.13%	8.57%	0.40	1.80%	6.06%	102.82	2.50%	6.47%
	ℓ	1000	–	–	0.01	1.38%	24.53%	0.20	0.39%	2.33%	25.03	0.11%	0.54%
	ℓ	2500	–	–	0.01	2.59%	21.19%	0.15	0.52%	3.61%	22.85	0.14%	0.96%
	ℓ^{pert}	1000	100	–	0.01	1.51%	23.61%	0.25	0.66%	4.27%	31.95	0.33%	2.23%
	ℓ^{pert}	2500	100	–	0.01	1.25%	19.54%	0.18	0.59%	3.46%	15.91	0.41%	2.00%
	ℓ	1000	–	LS	0.01	10.19%	54.67%	0.07	10.64%	51.23%	2.04	10.70%	46.99%
	ℓ	2500	–	LS	0.01	10.31%	55.70%	0.07	10.24%	50.63%	2.20	9.93%	43.87%
	ℓ^{pert}	1000	100	LS	0.01	10.28%	55.22%	0.07	10.64%	49.59%	2.47	10.70%	46.04%
	ℓ^{pert}	2500	100	LS	0.01	10.01%	53.64%	0.07	10.63%	49.27%	2.45	10.90%	46.36%
pert RDI \circ LS	FYL	–	–	–	0.36	0.02%	0.46%	2.58	0.02%	0.24%	208.72	0.02%	0.16%
	ℓ	1000	–	–	0.48	0.05%	1.53%	2.49	0.02%	0.62%	50.98	0.00%	0.10%
	ℓ	2500	–	–	0.50	0.09%	1.71%	2.44	0.02%	0.34%	45.37	0.00%	0.10%
	ℓ^{pert}	1000	100	–	0.49	0.05%	1.37%	2.61	0.02%	0.66%	65.53	0.00%	0.11%
	ℓ^{pert}	2500	100	–	0.48	0.05%	1.37%	2.54	0.02%	0.61%	40.28	0.00%	0.10%
	ℓ	1000	–	LS	0.48	0.03%	1.12%	2.34	0.01%	0.27%	14.06	0.00%	0.02%
	ℓ	2500	–	LS	0.49	0.03%	1.12%	2.32	0.00%	0.14%	14.19	0.00%	0.01%
	ℓ^{pert}	1000	100	LS	0.48	0.05%	2.26%	2.36	0.01%	0.27%	14.44	0.00%	0.05%
	ℓ^{pert}	2500	100	LS	0.49	0.05%	2.26%	2.39	0.01%	0.27%	14.59	0.00%	0.05%

T^{avg} is given in seconds.

Table 8: Influence of instances size on the performance of our solution algorithms with different \mathbf{w} on the $1|r_j|\sum_j C_j$ test set.

as a regret minimization problems. This widens the potential applications of such methods since it removes the need of an alternative algorithm for the hard problem to build the training set. Furthermore, even when such an algorithm exits, it may not be able to handle large instances. The learning by experience approach can therefore use larger instances in its training set, and can take into account the effect of potential post-processings. These two ingredients enable to scale better on large instances. Finally, we have shown that, if an approximation algorithm can be encoded in the pipeline with a given parametrization, then the parametrization learned by experience retains the approximation guarantee while giving a more efficient algorithm in practice.

Future contributions may focus on providing richer statistical models in the neural network, which would require to adapt the learning algorithm. Furthermore, the approximation ratio guarantee could be extended to more general settings.

Acknowledgements

I am grateful to Yohann de Castro and Julien Reynier for their help on Section 4, and to Vincent T'Kindt for his help on the scheduling problem.

References

- Brandon Amos and J. Zico Kolter. OptNet: Differentiable Optimization as a Layer in Neural Networks. In *Proceedings of the 34th International Conference on Machine Learning*, pages 136–145. PMLR. URL <https://proceedings.mlr.press/v70/amos17a.html>.
- Yoshua Bengio, Andrea Lodi, and Antoine Prouvost. Machine learning for combinatorial optimization: A methodological tour d’horizon. *European Journal of Operational Research*, 290(2):405–421, April 2021. doi: 10.1016/j.ejor.2020.07.063.
- Quentin Berthet, Mathieu Blondel, Olivier Teboul, Marco Cuturi, Jean-Philippe Vert, and Francis R. Bach. Learning with differentiable perturbed optimizers. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- Mathieu Blondel, André F. T. Martins, and Vlad Niculae. Learning with Fenchel-Young losses. 21(35):1–69. ISSN 1533-7928. URL <http://jmlr.org/papers/v21/19-021.html>.
- Mathieu Blondel, Quentin Berthet, Marco Cuturi, Roy Frostig, Stephan Hoyer, Felipe Llinares-López, Fabian Pedregosa, and Jean-Philippe Vert. Efficient and Modular Implicit Differentiation. In *Advances in Neural Information Processing Systems*, October 2022.
- Olivier Bousquet, Stéphane Boucheron, and Gábor Lugosi. Introduction to Statistical Learning Theory. In *Advanced Lectures on Machine Learning*, volume 3176, pages 169–207. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004. ISBN 978-3-540-23122-6 978-3-540-28650-9. doi: 10.1007/978-3-540-28650-9_8.
- Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.
- Guillaume Dalle, Léo Baty, Louis Bouvier, and Axel Parmentier. Learning with Combinatorial Optimization Layers: A Probabilistic Approach, July 2022.
- F. Della Croce and V. T’kindt. A recovering beam search algorithm for the one-machine dynamic total completion time scheduling problem. *Journal of the Operational Research Society*, 53:1275–1280, 2002.
- Adam N. Elmachtoub and Paul Grigas. Smart “Predict, then Optimize”. *Management Science*, March 2021.
- Bruno Escoffier, Laurent Gourvès, Jérôme Monnot, and Olivier Spanjaard. Two-stage stochastic matching and spanning tree problems: Polynomial instances and approximation. *European Journal of Operational Research*, 205(1):19–30, August 2010.

- Steven G. Johnson. The NLOpt nonlinear-optimization package. URL <http://github.com/stevengj/nlopt>. Accessed on 2021-07-04.
- D. R. Jones, C. D. Perttunen, and B. E. Stuckman. Lipschitzian optimization without the Lipschitz constant. *Journal of Optimization Theory and Applications*, 79(1):157–181, October 1993. doi: 10.1007/BF00941892.
- James Kotary, Ferdinando Fioretto, Pascal Van Hentenryck, and Bryan Wilder. End-to-End Constrained Optimization Learning: A Survey. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, pages 4475–4482. International Joint Conferences on Artificial Intelligence Organization, 2021. ISBN 978-0-9992411-9-6. doi: 10.24963/ijcai.2021/610. URL <https://www.ijcai.org/proceedings/2021/610>.
- Jayanta Mandi, Peter J Stuckey, Tias Guns, et al. Smart predict-and-optimize for hard combinatorial optimization problems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 1603–1610, 2020.
- Ruben Martinez-Cantin. BayesOpt: A Bayesian Optimization Library for Nonlinear Optimization, Experimental Design and Bandits. page 5.
- Sebastian Nowozin. Structured Learning and Prediction in Computer Vision. *Foundations and Trends® in Computer Graphics and Vision*, 6(3-4):185–365, 2010. doi: 10.1561/06000000033.
- Axel Parmentier. Learning to Approximate Industrial Problems by Operations Research Classic Problems. *Operations Research*, April 2021. doi: 10.1287/opre.2020.2094.
- Axel Parmentier and Vincent T’Kindt. Learning to solve the single machine scheduling problem with release times and sum of completion times. *arXiv:2101.01082 [cs, math]*, January 2021.
- Marin Vlastelica, Anselm Paulus, Vit Musil, Georg Martius, and Michal Rolinek. Differentiation of Blackbox Combinatorial Solvers. URL <https://openreview.net/forum?id=BkevoJSYPB>.
- Michael M. Wolf. Mathematical Foundations of Supervised Learning. https://www-m5.ma.tum.de/foswiki/pub/M5/Allgemeines/MA4801_2018S/ML_notes_main.pdf, 2018.

A Proof of Theorem 2

A.1 Background on Rademacher complexity and metric entropy method

This section introduces some classical tools of statistical learning theory [Bousquet et al., 2004]. The lecture notes of [Wolf, 2018] contain detailed proofs.

We place ourselves in the setting of Section 4.1. Let \mathcal{F} be the family of functions $\{\xi \mapsto \ell(\xi, \mathbf{w}) : \mathbf{w} \in \mathbf{W}\}$. The *Rademacher complexity* of \mathcal{F} is

$$\mathcal{R}_n(\mathcal{F}) = \mathbb{E}_{\xi_i, \sigma_i} \left[\sup_{\mathbf{w} \in \mathbf{W}} \frac{1}{n} \sum_{i=1}^n \sigma_i \ell(\xi_i, \mathbf{w}) \right]$$

where the σ_i are i.i.d. Rademacher variables, i.e., variables equal to 1 with probability 1/2, and to -1 otherwise. The following well-known result bounds the excess risk based on the Rademacher complexity.

Proposition 7. *With probability at least $1 - \delta$, we have*

$$L(\hat{\mathbf{w}}_n) - L(\mathbf{w}^*) \leq 4\mathcal{R}_n(\mathcal{F}) + \sqrt{\frac{2 \log(2/\delta)}{n}}.$$

The metric entropy method enables to bound the Rademacher complexity. The *empirical Rademacher complexity* of \mathcal{F} is obtained when we replace the expectation over ξ_i by its values for the training set used ξ_1, \dots, ξ_n .

$$\hat{\mathcal{R}}_n(\mathcal{F}) = \mathbb{E} \left[\sup_{\mathbf{w} \in \mathbf{W}} \frac{1}{n} \sum_{i=1}^n \sigma_i \ell(\xi_i, \mathbf{w}) \mid \xi_1, \dots, \xi_n \right]$$

and we have $\mathcal{R}_n(\mathcal{F}) = \mathbb{E}[\hat{\mathcal{R}}_n(\mathcal{F})]$.

Given n instances ξ_1, \dots, ξ_n and the corresponding distribution $\hat{\mu}_n$ on Ξ , the pseudometric $L_2(\hat{\mu}_n)$ on \mathcal{F} is the L_2 norm induced by $\hat{\mu}_n$ on \mathcal{F}

$$\|\ell(\cdot, \mathbf{w}) - \ell(\cdot, \mathbf{w}')\|_{2, \hat{\mu}_n} = \sqrt{\frac{1}{n} \sum_{i=1}^n (\ell(\xi_i, \mathbf{w}) - \ell(\xi_i, \mathbf{w}'))^2}$$

We denote by $B_{\varepsilon, L_2(\hat{\mu}_n)}(\ell(\cdot, \mathbf{w}))$ the ball of radius ε centered in $\ell(\cdot, \mathbf{w})$. The set covering number of \mathcal{F} with respect to $L_2(\hat{\mu}_n)$ is

$$N(\varepsilon, \mathcal{F}, L_2(\hat{\mu}_n)) = \min \left\{ m : \exists \{\mathbf{w}_1, \mathbf{w}_m\} \subseteq \mathbb{R}^d, \mathcal{F} \subseteq \bigcup_{j=1}^m B_{\varepsilon, L_2(\hat{\mu}_n)}(\ell(\cdot, \mathbf{w}_j)) \right\}.$$

The following result bounds the empirical Rademacher complexity from the covering number.

Proposition 8. (*Dudley's theorem*) *Let \mathcal{F} be a family of mapping from \mathcal{Z} to $[-1, 1]$, then*

$$\hat{\mathcal{R}}_n(\mathcal{F}) \leq 12 \int_0^\infty \sqrt{\frac{\log N(\varepsilon, \mathcal{F}, L_2(\hat{\mu}_n))}{n}} d\varepsilon$$

A.2 Proof of Theorem 2

The proof is as follows. We show that the Gaussian perturbation turns any bounded function in a Lipschitz function. Hence, the perturbed loss is Lipschitz. This implies an upper bound on the covering number, and Dudley's theorem enables to conclude.

Let \mathbf{Z} be a centered standard Gaussian vector on \mathbb{R}^d . It is well known that

$$\mathbb{E}(\|\mathbf{Z}\|) \leq \sqrt{d}. \quad (23)$$

Indeed, applying $u \leq (1+u^2)/2$ with $u = \sqrt{\frac{1}{d} \sum Z_i^2}$ gives $\frac{1}{\sqrt{d}}\|\mathbf{Z}\| \leq \frac{1}{2}(1 + \frac{1}{d} \sum_{i=1}^d Z_i^2)$. Taking the expectation and using $\mathbb{E}(Z_i^2) = 1$ gives (23).

Lemma 9. *Let $g : \mathbb{R}^d \rightarrow [0, 1]$ be an integrable function, \mathbf{Z} a standard normal random vector on \mathbb{R}^d , $\sigma > 0$ a positive real number, and $G(\mathbf{w}) = \mathbb{E}g(\mathbf{w} + \sigma\mathbf{Z})$. Then $\mathbf{w} \mapsto G(\mathbf{w})$ is $\frac{\sqrt{d}}{\sigma}$ -Lipchitz.*

Proof. Let h be the density of $\tilde{\mathbf{Z}} = \sigma\mathbf{Z}$. We have

$$G(\mathbf{w}) = \int h(\mathbf{z})g(\mathbf{z} + \mathbf{w}) = \int h(\mathbf{z} - \mathbf{w})g(\mathbf{z})$$

By dominated convergence, we have

$$\nabla G(\mathbf{w}) = - \int \nabla h(\mathbf{z} - \mathbf{w})g(\mathbf{z}) = - \int \nabla h(\mathbf{z})g(\mathbf{z} + \mathbf{w})$$

From there, using the facts that $|g(\mathbf{w})| \leq 1$ and \mathbf{Z} is a standard Gaussian, we get

$$\|\nabla G(\mathbf{w})\| \leq \int \|\nabla h(\mathbf{z})\| = \int \left\| \frac{\mathbf{z}}{\sigma^2(\sqrt{2\pi}\sigma)^n} e^{-\frac{\|\mathbf{z}\|^2}{2\sigma^2}} \right\| = \frac{\mathbb{E}(\|\tilde{\mathbf{Z}}\|)}{\sigma^2} = \frac{1}{\sigma}\mathbb{E}(\|\mathbf{Z}\|) \leq \frac{\sqrt{d}}{\sigma}$$

which gives the result. \square

Given an arbitrary element ξ in Ξ , Lemma 9 applied with $g = \ell(\xi, \cdot)$ gives

$$|\ell(\xi, \mathbf{w}) - \ell(\xi, \mathbf{w}')| \leq \frac{\sqrt{d}}{\sigma} \|\mathbf{w} - \mathbf{w}'\|_2$$

Hence

$$\|\ell(\cdot, \mathbf{w}) - \ell(\cdot, \mathbf{w}')\|_{2, \hat{\mu}_n} = \sqrt{\frac{1}{n} \sum_{i=1}^n (\ell(\xi_i, \mathbf{w}) - \ell(\xi_i, \mathbf{w}'))^2} \leq \frac{\sqrt{d}}{\sigma} \|\mathbf{w} - \mathbf{w}'\|.$$

As a consequence, if $\mathbf{w}_1, \dots, \mathbf{w}_m$ is an $\frac{\varepsilon\sigma}{\sqrt{d}}$ covering of \mathbf{W} endowed with the Euclidean norm, then $\ell(\cdot, \mathbf{w}_1), \dots, \ell(\cdot, \mathbf{w}_m)$ is an ε covering of \mathcal{F} . Hence, if \mathbf{W} is contained in the Euclidean ball of radius M , we get

$$N(\varepsilon, \mathcal{F}, L_2(\hat{\mu}_n)) \leq N\left(\frac{\varepsilon\sigma}{\sqrt{d}}, \mathbf{W} = \mathbb{B}^d(M), \|\cdot\|_2\right) \leq \left(\frac{M\sqrt{d}}{\varepsilon\sigma}\right)^d$$

for $\varepsilon \leq \frac{M\sqrt{d}}{\sigma}$ and $N(\varepsilon, \mathcal{F}, L_2(\hat{\mu}_n)) = 1$ otherwise. And we obtain

$$\log N(\varepsilon, \mathcal{F}, L_2(\hat{\mu}_n)) \leq d(\log(M\sqrt{d}/\sigma) - \log \varepsilon)$$

for $\varepsilon \leq \frac{M\sqrt{d}}{\sigma}$ and $\log N(\varepsilon, \mathcal{F}, L_2(\hat{\mu}_n)) = 0$ otherwise.

Proposition 8 then gives

$$\hat{R}_n(\mathcal{F}) \leq 12 \int_0^{\frac{M\sqrt{d}}{\sigma}} \sqrt{d \frac{\log(\frac{M\sqrt{d}}{\sigma}) - \log \varepsilon}{n}} d\varepsilon = 12 \sqrt{\frac{d}{n}} \int_0^{\frac{M\sqrt{d}}{\sigma}} \sqrt{-\log\left(\frac{\varepsilon}{M\sqrt{d}/\sigma}\right)} d\varepsilon = \frac{C}{4} \frac{Md}{\sigma\sqrt{n}}$$

with $C = 48 \int_0^1 \sqrt{-\log x} dx$. Remark that the bound on $\hat{R}_n(\mathcal{F})$ we obtain does not depend on the sample ξ_1, \dots, ξ_n , and is therefore also valid for $R_n(\mathcal{F}) = \mathbb{E}(\hat{R}_n(\mathcal{F}))$. Proposition 7 then gives Theorem 2.