

Learning to solve the single machine scheduling problem with release times and sum of completion times

Axel Parmentier^a, Vincent T'Kindt^{b,*}

^a*CERMICS, Ecole des Ponts, Marne-la-Vallée, France*

^b*University of Tours,*

LIFAT (EA 6300), ERL CNRS ROOT 7002, Tours, France

Abstract

In this paper, we focus on the solution of a hard single machine scheduling problem by new heuristic algorithms embedding techniques from machine learning field and scheduling theory. These heuristics transform an instance of the hard problem into an instance of a simpler one solved to optimality. The obtained schedule is then transposed to the original problem. Computational experiments show that they are competitive with state-of-the-art heuristics, notably on large instances.

Keywords: Scheduling, Single machine, Structured learning, Local search

1. Introduction

Consider the problem where n jobs have to be scheduled on a single machine. Each job j is defined by a *processing time* p_j and a release date r_j so that, in a give schedule, no job j can start before its release date. The machine can only process one job at a time and preemption is not allowed. The goal is to find a schedule s (permutation) that minimizes the total completion time $\sum_j C_j(s)$ with $C_j(s)$ the completion time of job j in schedule s . If $s = (j_1, \dots, j_n)$, then

$$C_{j_1}(s) = r_1 + p_1 \quad \text{and} \quad C_{j_k}(s) = \max(C_{j_{k-1}}(s), r_{j_k}) + p_{j_k} \quad \text{for } k > 1.$$

When there is no ambiguity, we omit the reference to schedule s when referring to completion times. Following the standard three-field notation in scheduling theory, this problem is referred to as $1|r_j|\sum_j C_j$ and is strongly \mathcal{NP} -hard [1]. When there is no release dates, the corresponding $1||\sum_j C_j$ problem can be solved in $O(n \log(n))$ time by means of the SPT rule (shortest processing times first). Again, when preemption is allowed, the corresponding $1|r_j, pmtn|\sum_j C_j$ problem can be solved $O(n \log(n))$ time by means of the SRPT rule (shortest remaining processing times first) [2].

*Corresponding author

Email addresses: axel.parmentier@enpc.fr (Axel Parmentier), tkindt@univ-tours.fr (Vincent T'Kindt)

The $1|r_j|\sum_j C_j$ problem is a challenging problem which has been studied for a long time. The two most competitive exact algorithms are the branch-and-memorize algorithm in [3] and the dynamic programming based algorithm in [4], which are able to solve instances with a hundred of jobs. Heuristic algorithms can be used to compute good solutions in a reasonable amount of time. Along the years, numerous heuristic algorithms have been proposed. We cite the RDI local search based on the APRTF greedy rule proposed in [5] and which requires $O(n^4 \log(n))$ time. The RBS heuristic (Recovering Beam Search) developed in [6] is a truncated search tree approach that has been the most performing heuristic for a decade. The RBS heuristic requires $O(w n^3 \log(n))$ time with w the beam width parametrizing the heuristic: notice that in [6] the case $w = 1$ is considered. To the best of our knowledge, the state-of-the-art heuristic is a matheuristic proposed in [7] which provides solutions very close to the optimal ones but at the price of a large CPU time requirement. All these heuristics have been milestones in the history of the $1|r_j|\sum_j C_j$ problem and will be considered as competitors for the learning based heuristics developed in this paper.




Figure 1: ML to approximate hard problems by well-solved ones

The use of machine learning (ML) techniques within operations research (OR) algorithms is a recent but active and promising research area [8]. To the best of our knowledge, very few contributions of this kind have considered scheduling problems [9, 10, 11]. In these works, ML is used to guide the solution process, *i.e.*, the proposed OR heuristic. In this paper, we elaborate on an original approach recently introduced in [12] and illustrated in Figure 1. A ML predictor φ_θ , which we call the encoding algorithm, is used to convert an instance Γ^h of a *hard* optimization problem into an instance Γ^e of an *easy* one, *i.e.*, a problem for which a practically efficient algorithm \mathcal{A}^e exists. Then, algorithm \mathcal{A}^e is used to compute an optimal solution s^e of Γ^e . And finally, a decoding algorithm is used to rebuild from s^e a solution s^h to Γ^h . In this paper, the hard problem is the $1|r_j|\sum_j C_j$ problem and the easy problem is the $1||\sum_j C_j$ problem. The processing time \hat{p}_j of job j in Γ^e is not equal to its processing time p_j in Γ^h , but to a linear combination of features computed from Γ^h . The easy problem algorithm \mathcal{A}^e is the SPT rule, and several different decoding algorithms are proposed to obtain a schedule for the $1|r_j|\sum_j C_j$ problem.

The main challenge to make such an approach work is to build an encoding algorithm φ_θ such that the optimal solution of the instance Γ^e leads to a good solution Γ^h after decoding. As usual in ML, we first define an appro-

appropriate family of predictors $(\varphi_{\theta})_{\theta}$, and then seek (learn) the best parameter θ . Following [12], we formulate the choice of θ as a structured learning problem. Supervised learning is the branch of machine learning that is interested in constructing an approximation of an unknown function $f : x \mapsto y$ from a training set $(x_1, y_1), \dots, (x_m, y_m)$, and structured learning [13] is the specific branch of supervised learning where the output y of f belongs to a structured and combinatorially large set \mathcal{Y} . In our case, \mathcal{Y} is the set of schedules, i.e., the set of permutations of $[n]$. Structured learning algorithms on the permutation group have been thoroughly studied in the literature on rankings [14], but with applications such as document retrieval, question answering or online advertising that are quite far from scheduling. If the traditional learning approaches of this literature can in theory be applied in our context, we do not use them because the loss functions they use to evaluate if a ranking is a good approximation of another are not good evaluations of the quality of a $1|r_j|\sum_j C_j$ schedule. We instead propose a novel approach based on Fenchel-Young [15] loss functions and inspired from the work of Berthet et al. [16]. This new approach is generic and can be applied to any hard problem for which the set of solutions of Γ^e does not depend on φ_{θ} .

The rest of the paper is organized as follows. Section 2 introduces the learning based heuristics developed to solve the $1|r_j|\sum_j C_j$ problem. The proposed structured learning approach is presented in Section 3. Extensive computational experiments showing the efficiency of the proposed heuristics are given in Section 4. Finally, Section 5 concludes the paper and provides further research directions.

2. Learning based heuristics for the $1|r_j|\sum_j C_j$ problem

To solve the $1|r_j|\sum_j C_j$ problem we propose three heuristic approaches. The general scheme of these heuristics is the one we already described on Figure 1: We first use a predictor φ_{θ} which transforms an instance $\Gamma^h = (n, (p_j, r_j)_{j \in [n]})$ into an instance $\Gamma^e = (n, (\hat{p}_j)_{j \in [n]})$ of the $1||\sum_j C_j$ problem. Then, the SPT rule solves the latter in polynomial time. And finally, we decode from the solution of Γ^e a solution of Γ^h . Consequently, the role of the predictor is mainly to define processing times \hat{p}_j for the $1||\sum_j C_j$ problem so that the schedule produced by the SPT rule is a good schedule when reintroducing the release dates r_j and the original processing times p_j . Let s^e be the schedule given by the SPT rule on Γ^e . Our three heuristics share the same predictor φ_{θ} , which is described in Section 2.1. The two first heuristics we propose are direct applications of the above principle, with different decoding algorithm and are developed in Section 2.2. The third heuristic is more elaborated and is developed in Section 2.3.

2.1. Encoding algorithm

Let us now focus on how the predictor φ_{θ} is built. A natural way to define the mapping φ_{θ} that maps Γ^h to an instance Γ^e is through a features vector

$\phi : (j, \Gamma^h) \mapsto \phi(j; \Gamma^h)$ that associates to each job j in Γ^h a vector of features $\phi(j; \Gamma^h)$. Each component of this vector of features provides information on job j . The choice of the features is a crucial point to make the approach efficient and it is discussed in Section 4. We suppose $\phi(i; \Gamma^h)$ to be in \mathbb{R}^d , with d the number of used features. Given a vector $\theta \in \mathbb{R}^d$ of parameters, we define the instance $\Gamma^e = \varphi_\theta(\Gamma^h)$ as the instance of $1||\sum_j C_j$ with n jobs whose processing times \hat{p}_j are defined by

$$\hat{p}_j = \langle \theta | \phi(j; \Gamma^h) \rangle \quad \forall j \in \{1, \dots, n\}.$$

In other words, the easy problem on instance Γ^h is

$$\min_s \langle \theta | \phi(s; \Gamma^h) \rangle \quad \text{where} \quad \phi(s; \Gamma^h) = \sum_{i=1}^n (n - i + 1) \phi(j_i; \Gamma^h) \quad (1)$$

and $s = (j_1, \dots, j_n)$. Note that depending on the choice of $\theta \in \mathbb{R}^d$, job processing times \hat{p}_j can be negative, which is not a problem since the SPT rule still applies in that case. The aim of the structured learning problem described in Section 3 is to choose $\theta \in \mathbb{R}^d$ in such a way that the optimal solution of Γ^e is a good solution for Γ^h .

2.2. Decoding algorithm: Heuristics PMLH and IMLH

Solutions of both $1|r_j|\sum_j C_j$ and $1||\sum_j C_j$ are schedules, i.e., permutations of $[n]$. Hence, a solution of $1||\sum_j C_j$ can be seen as a solution of $1|r_j|\sum_j C_j$, albeit with different completion times due to the release dates. The first heuristic we propose, denoted by PMLH (*Pure Machine Learning Heuristic*) and illustrated on Figure 2 therefore only returns the schedule s^e as a solution of Γ^h . The decoding algorithms only determine job starting times obtained when taking the release dates r_j and processing times p_j . It follows that heuristic PMLH requires $O(n \log(n))$ time.




Figure 2: Heuristic PMLH general scheme.

It may happen that, in the schedule built from s^e , two consecutive jobs j and k , with $j \rightarrow k$, are such that $p_j > p_k$. Let t be the starting time of j in that schedule. Then, if $r_j, r_k \leq t$ the schedule is suboptimal and swapping j and k leads to a better schedule. In the second heuristic we propose, denoted by IMLH (*Improved Machine Learning Heuristic*) and illustrated on Figure 3, we introduce in the decoding phase a fast local search LS to repair the above

problematic cases, if any. This local search is given in Algorithm 1 and it has a $O(n^2)$ worst-case time complexity. To improve the schedule s^i obtained from LS, we finally apply the RDI procedure described in [5]. More precisely, we apply the version of RDI where, each time a rescheduling of jobs is necessary, we use the SPT rule on the \hat{p}_j 's: always select among the available jobs at time t the one with smallest \hat{p}_j . The RDI heuristic requires $O(n^4 \log(n))$ time in the worst case. It follows that heuristic IMLH requires $O(n^4 \log(n))$ time in the worst case.




Figure 3: Heuristic IMLH general scheme.

Algorithm 1 The LS repairing heuristic

```

1: Input: An instance  $\Gamma^h$  and a schedule  $s^e$ .
2: Output: A schedule  $s^i$ .
3:  $t = 0, \ell = 0, s^i = \emptyset$ .
4:  $C_j = 0, \forall j = 1..n.$   $C_j$ : job completion times in  $s^i$ 
5: while ( $\ell < (n - 1)$ ) do
6:   if ( $t < r_{s^e[\ell]}$ ) then  $s^e[\ell]$  refers to the job in position  $\ell$  in  $s^e$ 
7:      $t = r_{s^e[\ell]}$ .
8:   end if
9:   if ( $t \geq r_{s^e[\ell+1]}$  and  $p_{s^e[\ell]} > p_{s^e[\ell+1]}$ ) then
10:     Swap  $s^e[\ell]$  and  $s^e[\ell + 1]$ .
11:     if ( $\ell = 0$ ) then
12:        $t = 0$ .
13:     end if
14:     if ( $\ell = 1$ ) then
15:        $t = 0, \ell = \ell - 1$ .
16:     end if
17:     if ( $\ell > 1$ ) then
18:        $t = C_{s^i[\ell-2]}, \ell = \ell - 1$ .
19:     end if
20:   else
21:      $t = t + p_{s^e[\ell]}, s^i[\ell] = s^e[\ell], C_{s^i[\ell]} = t, \ell = \ell + 1$ .
22:   end if
23: end while
24: return  $s^i$ .

```




Figure 4: Solution approach using several approximations

2.3. Decoding algorithm: Heuristic *itMLH*

Heuristic *itMLH* (iterative MLH) is an extension of heuristic *IMLH* since m vectors $\theta_k \in \mathbb{R}^d$ are generated by applying perturbations on the vector θ used in *IMLH*. The global scheme of *itMLH* is given in Figure 4. In this approach, we choose to approximate Γ^h by m instances of the easy problem $\Gamma_1^e, \dots, \Gamma_m^e$. We suggest to use the instances:

$$\Gamma_i^e = \varphi_{\theta_k}(\Gamma^h) \quad (2)$$

with $\theta_k = \theta + z_k$ and z_k is a sample of a random variable \mathbf{Z} on \mathbb{R}^d , that can be chosen arbitrarily. The processing time $\hat{p}_{j,k}$ of job j in instance Γ_k^e is thus:

$$\hat{p}_{j,k} = \langle \theta_k | \phi(j; \Gamma^h) \rangle.$$

Note that if the results of *PMLH* and *IMLH* depend on the direction of θ but not on its norm, the result of *itMLH* depends on its direction and its norm because of perturbation \mathbf{Z} . When learning the parameter θ for *itMLH*, we must therefore take into account the distribution of the perturbation \mathbf{Z} that is used in *itMLH*. This is what we do in Section 3. In our numerical experiments, we use a \mathbf{Z} distributed according to the normal distribution $\mathcal{N}(0, I_d)$ where I_d is the identity matrix of dimension d .

The heuristic *itMLH* starts by running *IMLH* before generating iteratively the m instances Γ_k^e . To speed up the heuristic, at each iteration k , if the schedule s^{e_k} has already been obtained at a previous iteration $k' < k$, then neither *LS* nor *RDI* are applied and the heuristic starts iteration $(k + 1)$. Similarly, *RDI* is not applied whenever s_k^i has already been computed at a previous iteration. It follows that heuristic *itMLH* requires $O(mn^4 \log(n))$ time in the worst case. Finally, remark that *itMLH* can be easily parallelized.

3. Structured learning methodology

3.1. Background on structured learning

After a brief introduction to structured learning which follows [12], we introduce our own variant of the perturbed optimizers and the resulting Fenchel-Young loss structured learning problem of Berthet et al. [16].

3.1.1. Setting

Broadly speaking, supervised learning [13] aims at learning an unknown function

$$\begin{aligned} f : \mathcal{X} &\rightarrow \mathcal{Y} \\ x &\mapsto y \end{aligned}$$

from a training set $(x_1, y_1), \dots, (x_n, y_n)$ where y_i is a noisy observation of $f(x_i)$. Structured learning [13] is a branch of supervised learning which deals with problems such that, for any x in \mathcal{X} , the mapping $f(x)$ takes its value in a set $\mathcal{Y}(x)$ that is finite, combinatorially large, and structured. To predict the value $f(x)$ on a new instance x , we solve the following optimization problem:

$$y^* = \hat{f}_\theta(x) := \arg \max_{y \in \mathcal{Y}(x)} g_\theta(y, x) \quad (3)$$

where the statistical model \hat{f}_θ is defined through an auxiliary *evaluation function* g_θ . Problem (3) is called the *structured prediction problem*. Given a training set $(x_i, y_i)_{i \in [n]}$, the objective of the *structured learning problem* is to learn a parameter θ such that \hat{f}_θ is a good approximation of f . Several approaches have been considered in the structured learning literature to formulate the structured learning problem, notably based on the maximum likelihood estimators (MLE) or on the minimization of surrogate loss functions [13]. Contrarily to [12], we do not use a MLE approach, because computing the likelihood would require to evaluate a sum on all the elements of the permutation group, which is a difficult problem that would require an ad-hoc algorithm. We instead introduce a Fenchel-Young loss approach, which leads to a much easier learning problem.

In the rest of the paper, we assume to have a *feature map* $\phi : (x, y) \mapsto \phi(y; x)$ which associates to every y in $\mathcal{Y}(x)$ a vector $\phi(y; x)$ in \mathbb{R}^d of features describing the properties of y . And we restrict ourselves to linear models of the form:

$$g_\theta(y, x) := \langle \theta | \phi(y; x) \rangle,$$

which leads to structured prediction problems of the form

$$\max_{y \in \mathcal{Y}(x)} \langle \theta | \phi(y; x) \rangle. \quad (4)$$

Let $\mathcal{C}(x)$ be the convex hull of $\{\phi(y; x) : y \in \mathcal{Y}(x)\}$ and $\mathcal{S}(x)$ be its linear span. Remark that it is useless to consider features $\phi(y; x)$ that are affinely dependent for all $y \in \mathcal{Y}$. Indeed, if $\mathcal{S}(x) \subsetneq \mathbb{R}^d$, and we decompose θ into $\theta_1 + \theta_2$ where θ_1 is the orthogonal projection of θ on $\mathcal{S}(x)$ and θ_2 on its orthogonal complement, then the value of θ_2 does not impact the result of the structured prediction problem. And we can obtain the same prediction flexibility by removing some features so as to obtain a full dimensional $\mathcal{C}(x)$. This justifies that, without loss of generality, we assume in the rest of the paper that $\mathcal{C}(x)$ has a non-empty interior (for at least some x). We denote by η the vectors in $\mathcal{S}(x)$.

3.1.2. Perturbed prediction problem

As described in Section 2.3, we add to $\boldsymbol{\theta}$ a random perturbation \mathbf{Z} , with \mathbf{Z} a random variable on \mathbb{R}^d . Let be:

$$F(\boldsymbol{\theta}; x) := \mathbb{E}_{\mathbf{Z}} \left(\max_{y \in \mathcal{Y}(x)} \langle \boldsymbol{\theta} + \mathbf{Z} | \boldsymbol{\phi}(y; x) \rangle \right) = \mathbb{E}_{\mathbf{Z}} \left(\max_{\boldsymbol{\eta} \in \mathcal{C}(x)} \langle \boldsymbol{\theta} + \mathbf{Z} | \boldsymbol{\eta} \rangle \right) \quad \text{and,}$$

$$\boldsymbol{\phi}^*(\boldsymbol{\theta}; \mathbf{x}) := \mathbb{E}_{\mathbf{Z}} \left(\boldsymbol{\phi} \left(\arg \max_{y \in \mathcal{Y}(x)} \langle \boldsymbol{\theta} + \mathbf{Z} | \boldsymbol{\phi}(y; x) \rangle; x \right) \right) = \mathbb{E}_{\mathbf{Z}} \left(\arg \max_{\boldsymbol{\eta} \in \mathcal{C}(x)} \langle \boldsymbol{\theta} + \mathbf{Z} | \boldsymbol{\eta} \rangle \right).$$

We denote by $\Omega(\cdot; x)$ the Fenchel dual of $F(\cdot; x)$, *i.e.*,

$$\Omega(\boldsymbol{\eta}; x) := \arg \max_{\boldsymbol{\theta} \in \mathcal{C}(x)} \langle \boldsymbol{\eta} | \boldsymbol{\theta} \rangle - F(\boldsymbol{\eta}; x).$$

The proof of the following result is available in Appendix A.

Proposition 1. The mapping $F(\cdot; x) : \boldsymbol{\theta} \mapsto F(\boldsymbol{\theta}; x)$ is convex with subgradient $\boldsymbol{\phi}^*(\boldsymbol{\theta}; \mathbf{x})$, which by abuse of notations we denote by $\nabla_{\boldsymbol{\theta}} F(\boldsymbol{\theta}; x)$.

1. If Z has positive and differentiable density $d\mu(z) \propto \exp(-\nu(z))dz$, then $F(\cdot; x)$ is twice differentiable and the subgradient above is its gradient. If in addition $\mathcal{C}(x)$ has non-empty interior, then $F(\cdot; x)$ is strictly convex.
2. If Z has a sampled distribution $\frac{1}{m} \sum_{i=1}^m \delta_{\mathbf{z}_i}$, then $F(\cdot; x)$ is piecewise linear.

3.1.3. Perturbed model learning with Fenchel-Young losses

In the loss approach to supervised learning, we suppose to have a loss $\ell(\boldsymbol{\theta}; y, x)$ that evaluates how far the y predicted by (3) is from the true y , and we formulate the learning problem as

$$\min_{\boldsymbol{\theta}} \frac{1}{n} \sum_{i=1}^n \ell(\boldsymbol{\theta}; y_i, x_i).$$

An ideal loss $\ell(\boldsymbol{\theta}; y, x)$ for the structured learning problem would be a loss that is non-negative, easy to minimize and such that $\ell(\boldsymbol{\theta}; y, x) = 0$ if and only if the y^* realizing the maximum in (4) is such that $y^* = y$. But since predictions are done using (4) where y appears only through $\boldsymbol{\phi}(y; x)$, we cannot hope to distinguish y_1 and y_2 if $\boldsymbol{\phi}(y_1; x) = \boldsymbol{\phi}(y_2; x)$. Thus, all we can ask for is that $\ell(\boldsymbol{\theta}; y, x) = 0$ if and only if $\boldsymbol{\phi}(y; x) \in \arg \max_{\boldsymbol{\eta} \in \mathcal{C}(x)} \langle \boldsymbol{\theta} | \boldsymbol{\eta} \rangle$.

Given x and $y \in \mathcal{Y}(x)$, the *Fenchel-Young loss* $L(\boldsymbol{\theta}; y, x)$ is defined by:

$$L(\boldsymbol{\theta}; y, x) = F(\boldsymbol{\theta}; x) + \Omega(\boldsymbol{\phi}(y; x); x) - \langle \boldsymbol{\theta} | \boldsymbol{\phi}(y; x) \rangle.$$

It follows from Proposition 1 that the following vector, which by abuse of notations we denote by $\nabla_{\boldsymbol{\theta}} L(\boldsymbol{\theta}; y, x)$, is a subgradient of $L(\cdot; y, x)$.

$$\nabla_{\boldsymbol{\theta}} L(\boldsymbol{\theta}; y, x) = \boldsymbol{\phi}^*(\boldsymbol{\theta}; x) - \boldsymbol{\phi}(y; x) \tag{5}$$

Given \mathbf{z} in \mathbb{R}^d , we denote by $\delta_{\mathbf{z}}$ the Dirac distribution in \mathbf{z} . The proof of the following result is available in Appendix A.

Proposition 2. The Fenchel-Young loss $L(\cdot; y, x): \boldsymbol{\theta} \mapsto L(\boldsymbol{\theta}; y, x)$ is non-negative, convex in $\boldsymbol{\theta}$, and with value 0 if and only if $\boldsymbol{\phi}(y; x) \in \partial_{\boldsymbol{\theta}} F(\boldsymbol{\theta}; x)$, where $\partial_{\boldsymbol{\theta}} F(\cdot; x)$ is the subdifferential of $F(\cdot; x)$.

1. If Z has positive and differentiable density $d\mu(z) \propto \exp(-\nu(z))dz$ and $\mathcal{C}(x)$ has non-empty interior, then $L(\cdot; y, x)$ is strictly convex and $L(\boldsymbol{\theta}; y, x) = 0$ if and only if $\boldsymbol{\phi}(y; x) = \boldsymbol{\phi}^*(\boldsymbol{\theta}; x)$.
2. If Z has a sampled distribution $\frac{1}{m} \sum_{i=1}^m \delta_{\mathbf{z}_i}$, then $L(\boldsymbol{\theta}; y, x) = 0$ if and only if there exists $\boldsymbol{\eta}_1 \in \arg \max_{\boldsymbol{\eta} \in \mathcal{C}(x)} \langle \boldsymbol{\theta} + \mathbf{z}_1 | \boldsymbol{\eta} \rangle, \dots, \boldsymbol{\eta}_m \in \arg \max_{\boldsymbol{\eta} \in \mathcal{C}(x)} \langle \boldsymbol{\theta} + \mathbf{z}_m | \boldsymbol{\eta} \rangle$ such that $\boldsymbol{\phi}(y; x) = \frac{1}{m} \sum_{i=1}^m \boldsymbol{\eta}_i$.

Given a training set $(x_1, y_1), \dots, (x_n, y_n)$, we therefore formulate the learning problem as

$$\min_{\boldsymbol{\theta} \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n L(\boldsymbol{\theta}; y_i, x_i). \quad (6)$$

In practice, we use a normal distribution $\mathcal{N}(0, I_d)$ for \mathbf{Z} , and the first point of Proposition (2) ensures that (6) is a strictly convex optimization problem provided that $\mathcal{C}(x_i)$ has non-empty interior for at least one i , which is practically the case as we explained at the end of Section 3.1.1.

The practical difficulty of the problem (6) lies in the integrals that appear in F and its gradient, even if we omit the hard to compute constant terms $\Omega(y_i)$ in the objective function. Therefore, we replace it by its *sample average approximation* (SAA): we draw m samples $\mathbf{z}_1, \dots, \mathbf{z}_m$ of Z and replace μ by its sample average approximation $\tilde{\mu}$. The consequence is to take the expectations with respect to $\tilde{\mu}$ when computing F in the objective function and $\boldsymbol{\phi}_\varepsilon^*$ in its gradient (5). The second point of Proposition 2 then enables to interpret the $\boldsymbol{\theta}$ obtained. And the second point of Proposition 1 ensures that the SAA problem is a convex and piecewise-linear non-constrained optimization problem. Given that (5) ensures that subgradients can be computed by solving several problems of the form (4), we can therefore solve the SAA problem using a subgradient algorithm, or any other non-differentiable convex optimization method. Lewis and Overton [17] and others have observed that BFGS algorithms perform well in practice on many such problems, although there is no theoretical convergence guarantee. We therefore solve the problem with a BFGS algorithm. And we have observed a fast convergence to an optimal solution – the certificate being a null gradient. Alternative approaches could of course be used.

Remark 1. The reader familiar with Machine Learning will note that we do not add a regularization term in (6). Indeed, the perturbation \mathbf{Z} can be seen as a regularizer on the features space, as it is generally the case with Fenchel Young losses (see Blondel et al. [15] for more details).

Remark 2. Berthet et al. [16] solve the non-sampled learning problem using a stochastic gradient descent, which is required in their case because their $\boldsymbol{\phi}(y; x)$ is the output of a deep neural network that they fit using this algorithm. As

we do not have this practical constraint, we have chosen to minimize the SAA approximation with a BFGS algorithm because it is easy to implement and gives a fast convergence in practice.

3.2. Learning to approximate the $1|r_j|\sum_j C_j$ problem

We formulate the problem of learning the parameter θ of our scheme of Section 2 as a structured learning problem. In this case, the function f we approximate is the function:

$$f : \Gamma \mapsto s^*(\Gamma) \quad \text{where } s^*(\Gamma) \text{ is an optimal solution } \Gamma,$$

and Γ is an instance of the $1|r_j|\sum_j C_j$ problem. The structured prediction problem (4) then coincides with the easy problem (1), the only difference being that the maximum is replaced by a minimum. All we have to do is therefore to build a training set $(\Gamma_1, s_1), \dots, (\Gamma_n, s_n)$ of instances of the scheduling problem with their optimal solution and solve the learning problem of Section 3.1.3 to obtain a parameter θ^* . We can then use $-\theta^*$ as parameter θ in (1).

Remark 3. In their work on perturbed maximizers, Berthet et al. [16] perturb the feature vector ϕ instead of the parameter θ . For a reader that would compare the approaches, we underline that our ϕ correspond to their θ (see notably Section 5.2 in [16]). We have chosen to use this alternative perturbation because it leads to better numerical results on the problem we consider.

3.3. Broad applicability

In this paper, we use a structured learning approach to approximate instances of the hard problem $1|r_j|\sum_j C_j$ by easier instances of the $1|\sum_j C_j$ problem. As detailed in [12], the paradigm of using machine learning to approximate instances of a hard problem by instances of an easier one is quite generic, and can be used, for instance, to linearize problems. However, the structured learning algorithm used to learn θ in [12] is not generic: ad-hoc algorithms must be derived for new applications. A strong advantage of the approach proposed in paper is that it is completely generic. Indeed, the only things we need to apply it to new problems is:

1. a training set composed of hard problem instances with their optimal solutions (or good quality solutions),
2. a vector of features $\phi(y; x)$ – See Section 1.3 of [12] for details on how such vector of features can be built,
3. an algorithm solving the easy problem.

Indeed, an algorithm solving instances of the easy problem is all we need to compute $F(\theta; x)$ and its gradient. And such an algorithm is always available: the existence of such an algorithm is the reason why we want to approximate the hard problem by this specific easy problem. The structured learning approach based on Fenchel-Young losses detailed in this paper is therefore much more generic than the MLE approach, because it does not require additional algorithms.

4. Computational experiments

In this section we focus on the experiments done to build the predictor φ_{θ} and the comparisons of the learning based heuristics with respect to state-of-the-art algorithms. We first introduce how instances Γ^h of the $1|r_j|\sum_j C_j$ problem are randomly generated ([6]). For a given instance of n jobs, processing times p_j are drawn at random following the uniform distribution $[1; 100]$ and release times r_j are drawn at random following the uniform distribution $[1; 50.5 * n * \rho]$. 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 values of n and ρ , N instances are randomly generated leading for a fixed value of n to $10 * N$ instances. The values considered for n and N depend on the conducted experiments and are given in the next sections.

The algorithms considered in the experiments are the learning based heuristics PMLH, IMLH and itMLH with $m = 150$. To evaluate the contribution of the predictor, we also implemented a version of IMLH in which instead of using the predictor to build an initial solution, we randomly generate it. Next, LS local search is applied to repair local suboptimal jobs sequencing. This heuristic is denoted by RAND.

We also use the RDI/APRTF heuristic ([5]), referred to as RDIA, the RBS heuristic with beam width $w = 2$ ([6]), referred to as RBS, and the matheuristic ([7]), referred to as MATH. Besides, for instances with up to $n = 110$ jobs, the optimal solution is computed by means of the branch-and-memory algorithm in [3]. All heuristic algorithms but MATH have been coded in C++ language. The code of heuristic MATH has been kindly provided by Fabio Salassa ([7]). Testings have been done on a PC Intel XEON E5 with 8 cores of 2.20Ghz and 8Gb of RAM.

4.1. Features definition and selection

To build a predictor φ_{θ} , the vector of features $\phi(j; \Gamma^h)$ must be defined. We initially considered 66 features, and manually selected $d = 27$ of these features that lead to good prediction performances. This selection was done using a validation set distinct from the test set later used. The $\phi(j; \Gamma^h)$ vector of (normalized) features is given in Table 1. We denote by $[j]^X$ the position of job j in the sequence obtained by applying the sequencing rule $X \in \{SPT, SRT, SP+RT\}$. Rule SRT (resp. SP+RT) consists in sorting the jobs by increase value of the r_j 's (resp. $r_j + p_j$'s). Besides, features 15-17 and 23-27 are computed exploiting the SRPT algorithm ([2]) which solves the preemptive version of the $1|r_j|\sum_j C_j$ problem, thus enabling to get in polynomial time a lower bound to the original problem. In a preemptive schedule computed by SRPT, some jobs j can be interrupted at time r_k , while being processed, by the arrival of a job k such that $p_k < (p_j - \pi_j)$ with π_j the duration j has been processed before its first preemption by job k . In case job j is not preempted, we have $\pi_j = p_j$. Let $\#_j$ be the number of times job j is preempted by another job, and $\#_T$ be the total number of preemptions. Finally, $[j]^{SRPT}$ is the position of the completion time of job j in the SRPT schedule. For feature 16, p_k refers to the processing time of the job that follows the initial part of job j and of duration π_j .

For features 18-21, $Decile(X)$ refers to the the decile number of $X \in \{r_j, p_j\}$ when the X 's are sorted by increasing value. For features 24-27, set $BS_j^p = \{k \mid [k]^{SRPT} < [j]^{SRPT}, p_k < p_j\}$, set $BS_j^r = \{k \mid [k]^{SRPT} < [j]^{SRPT}, r_k < r_j\}$, set $BG_j^p = \{k \mid [k]^{SRPT} < [j]^{SRPT}, p_k > p_j\}$ and set $BG_j^r = \{k \mid [k]^{SRPT} < [j]^{SRPT}, r_k > r_j\}$.

| | Description | θ_k | σ_k | | Description | θ_k | σ_k |
|----|--|-------------|------------|----|--|------------|------------|
| 1 | $\frac{[j]^{SPT}}{n}$ | 4.05111 | 0.28865 | 15 | $\frac{p_j - \pi_j}{\sum_{i=1}^n p_i - \pi_i}$ | -38.67500 | 0.03549 |
| 2 | $\frac{[j]^{SKT}}{n}$ | -11.38040 | 0.28865 | 16 | $\frac{p_j - \pi_j}{p_k \sum_{i=1}^n p_i - \pi_i}$ | 8.72219 | 0.04948 |
| 3 | $\frac{[j]^{SP+RT}}{n}$ | -10.19020 | 0.28865 | 17 | $\frac{p_j - \pi_j}{p_j \sum_{i=1}^n p_i - \pi_i}$ | 40.31390 | 0.03392 |
| 4 | $\frac{r_j}{p_j} \times \frac{\sum_{i=1}^n p_i}{\sum_{i=1}^n r_i}$ | 1.89904 | 0.029879 | 18 | $Decile(r_j)$ | 0.84073 | 2.86997 |
| 5 | $\frac{p_j}{r_j} \times \frac{\sum_{i=1}^n r_i}{\sum_{i=1}^n p_i}$ | -21.23830 | 0.04374 | 19 | $\frac{r_j}{Decile(r_j)}$ | 206.65600 | 0.00482 |
| 6 | $\frac{r_j}{\sum_{i=1}^n r_i}$ | 5440.67000 | 0.00839 | 20 | $Decile(p_j)$ | 0.12019 | 2.86839 |
| 7 | $\frac{p_j}{\sum_{i=1}^n r_i}$ | 6467.43000 | 0.00077 | 21 | $\frac{p_j}{Decile(p_j)}$ | 87.99630 | 0.00473 |
| 8 | $\frac{r_j + p_j}{\sum_{i=1}^n r_i}$ | 5286.99000 | 0.00856 | 22 | $\frac{\#j}{\#T}$ | -31.23980 | 0.03438 |
| 9 | $\frac{r_j}{\sum_{i=1}^n p_i}$ | 47.20590 | 0.59882 | 23 | $\frac{[j]^{SRPT}}{n}$ | 125.59400 | 0.28865 |
| 10 | $\frac{p_j}{\sum_{i=1}^n p_i}$ | -345.58400 | 0.00835 | 24 | $\frac{ BS_j^p }{\sum_{i=1}^n BS_i^p }$ | 162.31700 | 0.01218 |
| 11 | $\frac{r_j + p_j}{\sum_{i=1}^n p_i}$ | 47.14900 | 0.59890 | 25 | $\frac{ BS_j^r }{\sum_{i=1}^n BS_i^r }$ | 429.34900 | 0.00878 |
| 12 | $\frac{r_j}{\sum_{i=1}^n r_i + p_i}$ | -6733.09000 | 0.00802 | 26 | $\frac{ BG_j^p }{\sum_{i=1}^n BG_i^p }$ | 34.17510 | 0.01473 |
| 13 | $\frac{p_j}{\sum_{i=1}^n r_i + p_i}$ | -3879.31000 | 0.00066 | 27 | $\frac{ BG_j^r }{\sum_{i=1}^n BG_i^r }$ | 39.68920 | 0.03044 |
| 14 | $\frac{r_j + p_j}{\sum_{i=1}^n r_i + p_i}$ | -6555.45000 | 0.00814 | | | | |

Table 1: List of features of a given job j , and their value θ_k in the learned θ .

4.2. Learning algorithm

A training database has been built by randomly generating instances Γ^h with $n \in \{50, 70, 90, 110\}$ and $N = 100$, thus leading to 4000 entries in the database. Each entry is defined by:

$$(n, \sum_j C_j^*, r_1, p_1, \phi(1; \Gamma^h), \dots, r_n, p_n, \phi(n; \Gamma^h)),$$

with $\sum_j C_j^*$ the optimal solution value for Γ^h . After learning, we obtain the normalized θ vector, as well as their standard deviations σ , given in Table 1. Generating the database took 2 hours of computing time. We then build the the sample average approximation of the learning problem of Section 3.1.3 by drawing 100 samples of \mathbf{Z} . Evaluating the objective of the learning problem thus requires to solve 4000×100 instances of the easy problem with the SPT rule. We solve the learning problem with the BFGS implementation of [18]. The BFGS algorithms converges after 91 iterations, which took a total of 1 hour and 31 minutes.

4.3. Comparisons of the heuristics

Comparison to optimal solutions.

We first focus on the comparison of the heuristics with the optimal solution on instances Γ^h with $n \in \{50, 60, 70, 80, 90, 100, 110\}$ and $N = 30$. None of the randomly generated instances corresponds to some instances generated when building the training database. Table 2 presents the obtained results. For each problem size n and each heuristic $H \in \{\text{RAND}, \text{PMLH}, \text{IMLH}, \text{iTMLH}, \text{RDIA}, \text{RBS}, \text{MATH}\}$, we compute several statistics. Column δ_{avg} (resp. δ_{max}) is the average (resp. maximum) deviation of heuristic H to the optimal solution of the N instances of size n . For a given instance Γ^h , the deviation δ is computed as follows:

$$\delta = 100.00 \times \frac{\sum_j C_j(H) - \sum_j C_j^*}{\sum_j C_j^*},$$

with $\sum_j C_j(H)$ the value of the solution returned by H and $\sum_j C_j^*$ the optimal solution value. Column $\#Opt$ is the percentage of instances for which heuristic H finds an optimal solution. Column T_{avg} (resp. T_{max}) is the average (resp. maximum) CPU time in seconds used by heuristic H .

First, the results in Table 2 show that heuristic **MATH** strongly outperforms all other heuristics but at the price of a CPU time significantly higher. Heuristic **RBS** is slightly worse than **MATH** and with a running time significantly higher than for the other heuristics. These results are in line with those given in [7]. The poor performance of **RAND** shows that the quality of the prediction made by φ_θ is crucial for the performance of **IMLH**: the local search alone does not enable to find a good solution. Furthermore, even **PMLH**, which does not perform a local optimization after the prediction gives interesting results in terms of deviations to optimality. Besides, heuristic **IMLH** slightly outperforms heuristic **RDIA** both in terms of deviations and percentage of instances solved to optimality. Remind that the only difference between these two heuristics is the initial solution given to the RDI local search: **IMLH** uses a predictor and a fast local search, while **RDIA** uses a dedicated heuristic algorithm. Finally, we can remark that heuristic **iTMLH**, that improves upon **IMLH**, provides deviations to optimality slightly worse than those of **RBS** but in a reduced running time.

We can conclude from these experiments that:

1. Learning based heuristics are competitive with the state-of-the-art heuristics **RDIA** and **RBS**.
2. Learning based heuristics require a reduced running time with respect to **RBS** and **MATH**.
3. Heuristic **MATH** is a way far the most effective heuristic in terms of deviations to optimality.

Comparison to the best known solutions.

We now focus on the comparison of the heuristics with the best known solution on larger instances Γ^h with $n \in \{120, 140, 160, 180, 200, 300, 500, 1000, 1500, 2000\}$,

| n | RAND | | | | | PMLH | | | | | IMLH | | | | |
|-----|--------------------|--------------------|-------------|--------------|--------------|--------------------|--------------------|-------------|--------------|--------------|--------------------|--------------------|-------------|--------------|--------------|
| | $\delta_{avg}(\%)$ | $\delta_{max}(\%)$ | $\#Opt(\%)$ | $T_{avg}(s)$ | $T_{max}(s)$ | $\delta_{avg}(\%)$ | $\delta_{max}(\%)$ | $\#Opt(\%)$ | $T_{avg}(s)$ | $T_{max}(s)$ | $\delta_{avg}(\%)$ | $\delta_{max}(\%)$ | $\#Opt(\%)$ | $T_{avg}(s)$ | $T_{max}(s)$ |
| 50 | 77.779 | 134.602 | 0.00 | 0.00 | 0.00 | 1.491 | 6.643 | 0.00 | 0.00 | 0.00 | 0.208 | 2.371 | 18.33 | 0.01 | 1.00 |
| 60 | 80.699 | 150.660 | 0.00 | 0.00 | 0.00 | 1.212 | 5.231 | 0.00 | 0.01 | 1.0 | 0.181 | 3.720 | 18.00 | 0.01 | 1.00 |
| 70 | 84.460 | 139.698 | 0.00 | 0.00 | 0.00 | 1.066 | 5.533 | 0.00 | 0.00 | 0.00 | 0.171 | 1.784 | 11.67 | 0.00 | 1.00 |
| 80 | 86.597 | 147.243 | 0.00 | 0.01 | 1.00 | 0.994 | 4.857 | 0.00 | 0.00 | 1.0 | 0.157 | 2.647 | 7.00 | 0.01 | 1.00 |
| 90 | 89.537 | 158.000 | 0.00 | 0.00 | 0.00 | 0.973 | 4.414 | 0.00 | 0.00 | 0.00 | 0.128 | 1.481 | 10.67 | 0.01 | 1.00 |
| 100 | 90.700 | 144.174 | 0.00 | 0.00 | 1.00 | 0.919 | 4.276 | 0.00 | 0.01 | 1.00 | 0.118 | 1.578 | 6.67 | 0.01 | 1.00 |
| 110 | 91.666 | 145.854 | 0.00 | 0.00 | 0.00 | 0.903 | 3.848 | 0.00 | 0.01 | 1.00 | 0.103 | 1.209 | 8.33 | 0.01 | 1.00 |

| n | itMLH | | | | | RDIA | | | | | RBS | | | | |
|-----|--------------------|--------------------|-------------|--------------|--------------|--------------------|--------------------|-------------|--------------|--------------|--------------------|--------------------|-------------|--------------|--------------|
| | $\delta_{avg}(\%)$ | $\delta_{max}(\%)$ | $\#Opt(\%)$ | $T_{avg}(s)$ | $T_{max}(s)$ | $\delta_{avg}(\%)$ | $\delta_{max}(\%)$ | $\#Opt(\%)$ | $T_{avg}(s)$ | $T_{max}(s)$ | $\delta_{avg}(\%)$ | $\delta_{max}(\%)$ | $\#Opt(\%)$ | $T_{avg}(s)$ | $T_{max}(s)$ |
| 50 | 0.055 | 1.157 | 26.00 | 0.51 | 1.00 | 0.229 | 2.687 | 19.00 | 0.00 | 0.00 | 0.014 | 0.302 | 67.00 | 0.43 | 1.00 |
| 60 | 0.048 | 0.571 | 23.33 | 0.55 | 1.00 | 0.191 | 2.058 | 13.00 | 0.00 | 0.00 | 0.017 | 0.426 | 62.00 | 0.73 | 1.00 |
| 70 | 0.048 | 1.079 | 18.33 | 0.67 | 1.00 | 0.218 | 1.804 | 12.67 | 0.00 | 0.00 | 0.015 | 0.158 | 57.33 | 1.31 | 3.00 |
| 80 | 0.052 | 1.675 | 10.00 | 0.67 | 2.00 | 0.209 | 2.018 | 8.67 | 0.00 | 1.00 | 0.016 | 0.193 | 46.67 | 1.95 | 3.00 |
| 90 | 0.043 | 0.342 | 11.67 | 0.71 | 2.00 | 0.204 | 1.757 | 6.67 | 0.00 | 0.00 | 0.017 | 0.388 | 44.33 | 2.80 | 5.00 |
| 100 | 0.037 | 0.359 | 8.33 | 0.80 | 2.00 | 0.182 | 2.273 | 7.33 | 0.00 | 0.00 | 0.013 | 0.162 | 43.33 | 3.96 | 7.00 |
| 110 | 0.034 | 0.280 | 8.67 | 0.90 | 2.00 | 0.181 | 1.760 | 6.33 | 0.00 | 1.00 | 0.013 | 0.120 | 37.67 | 5.54 | 7.00 |

| n | MATH | | | | |
|-----|--------------------|--------------------|-------------|--------------|--------------|
| | $\delta_{avg}(\%)$ | $\delta_{max}(\%)$ | $\#Opt(\%)$ | $T_{avg}(s)$ | $T_{max}(s)$ |
| 50 | 0.000 | 0.042 | 98.33 | 4.82 | 30.00 |
| 60 | 0.002 | 0.149 | 98.00 | 7.44 | 34.00 |
| 70 | 0.001 | 0.067 | 95.67 | 10.53 | 46.00 |
| 80 | 0.003 | 0.151 | 92.67 | 11.53 | 59.00 |
| 90 | 0.004 | 0.377 | 90.33 | 14.00 | 81.00 |
| 100 | 0.001 | 0.038 | 92.00 | 16.88 | 67.00 |
| 110 | 0.001 | 0.031 | 90.33 | 20.85 | 68.00 |

Table 2: Comparison of the heuristics with optimal solutions

2500} and $N = 30$. For a given instance Γ^h , the best known solution is the best solution found by the heuristics. None of the randomly generated instances corresponds to some instances generated when building the training database. Also notice that a time limit of 180s is imposed in the experiments: as soon as for a given size n , the average running time of a heuristic exceeds 180s, this one is no longer run for higher values of n . The heuristic RAND is not included in these experiments.

Table 3 presents the obtained results. The meaning of the columns is the same than for Table 2, with the modification that, for a given instance Γ^h , the deviation δ is computed as follows:

$$\delta = 100.00 \times \frac{\sum_j C_j(H) - \sum_j C_j^{BNS}}{\sum_j C_j^{BNS}},$$

with $\sum_j C_j(H)$ the value of the solution returned by H and $\sum_j C_j^{BNS} = \min_{H \text{ is running}}(\sum_j C_j(H))$ the value of the best known solution.

Table 3 shows that heuristics MATH and RBS are unable to solve instances with more than $n = 300$ jobs within the time limit of 180s. Heuristic RDIA is able to solve instances with up to 2000 jobs but is slower than IMLH and itMLH. Regarding the deviations to the best known solutions, heuristic MATH remains the most effective one, with an average deviation almost equal to 0. Heuristic RBS is the second most effective heuristic but, again, a slow one not able to solve large instances. Heuristics IMLH and itMLH provide very good results in terms of deviations with reduced average running times.

We can conclude from these experiments that:

1. Learning based heuristics offer a very good trade-off between the quality of the computed solution and the running time required. They also show very low deviations to the best known solution.
2. Heuristic RDIA is outperformed by both IMLH and itMLH.
3. Heuristics MATH and RBS are enable to solve instances with more than 300 jobs within the time limit of 180s.

5. Conclusion

In summary, we introduce several new heuristics for the scheduling problem $1|r_j|\sum_j C_j$. These heuristics all rely on a structured learning predictor used to approximate an instance of $1|r_j|\sum_j C_j$ by an instance of $1||\sum_j C_j$. Our heuristics are competitive with RDI-APRTF, the state-of-the-art “fast” heuristic for $1|r_j|\sum_j C_j$ on instances with up to 500 jobs, and outperform it on larger instances. To the best of our knowledge, they are the first machine learning based algorithms to outperform state-of-the-art algorithms on a scheduling problem. And research directions include the extension of our approach to other scheduling problems.

| n | PMLH | | | | IMLH | | | | itMLH | | | |
|------|--------------------|--------------------|--------------|--------------|--------------------|--------------------|--------------|--------------|--------------------|--------------------|--------------|--------------|
| | $\delta_{avg}(\%)$ | $\delta_{max}(\%)$ | $T_{avg}(s)$ | $T_{max}(s)$ | $\delta_{avg}(\%)$ | $\delta_{max}(\%)$ | $T_{avg}(s)$ | $T_{max}(s)$ | $\delta_{avg}(\%)$ | $\delta_{max}(\%)$ | $T_{avg}(s)$ | $T_{max}(s)$ |
| 120 | 0.927 | 3.610 | 0.02 | 1.00 | 0.106 | 1.279 | 0.03 | 1.00 | 0.034 | 0.648 | 0.97 | 2.00 |
| 140 | 0.886 | 3.780 | 0.01 | 1.00 | 0.076 | 1.389 | 0.02 | 1.00 | 0.032 | 0.244 | 1.19 | 2.00 |
| 160 | 0.930 | 3.691 | 0.02 | 1.00 | 0.079 | 0.740 | 0.04 | 1.00 | 0.034 | 0.276 | 1.26 | 3.00 |
| 180 | 0.974 | 4.804 | 0.01 | 1.00 | 0.081 | 0.665 | 0.02 | 1.00 | 0.029 | 0.146 | 1.42 | 2.00 |
| 200 | 0.981 | 4.360 | 0.01 | 1.00 | 0.078 | 1.326 | 0.05 | 1.00 | 0.025 | 0.136 | 1.68 | 3.00 |
| 300 | 1.146 | 5.821 | 0.02 | 1.00 | 0.060 | 0.750 | 0.05 | 1.00 | 0.025 | 0.193 | 2.37 | 5.00 |
| 500 | 1.275 | 4.608 | 0.04 | 1.00 | 0.029 | 0.390 | 0.21 | 1.00 | 0.005 | 0.091 | 3.82 | 7.00 |
| 1000 | 1.453 | 4.885 | 0.07 | 1.00 | 0.024 | 0.234 | 2.04 | 9.00 | 0.007 | 0.069 | 9.85 | 22.00 |
| 1500 | 1.520 | 4.749 | 0.13 | 1.00 | 0.020 | 0.190 | 11.35 | 41.00 | 0.008 | 0.068 | 26.40 | 84.00 |
| 2000 | 1.555 | 5.069 | 0.13 | 1.00 | 0.017 | 0.116 | 40.37 | 146.00 | 0.009 | 0.070 | 71.43 | 300.00 |
| 2500 | 1.559 | 4.668 | 0.16 | 1.00 | 0.008 | 0.135 | 108.68 | 472.00 | 0.000 | 0.000 | 185.44 | 907.00 |

| n | RDIA | | | | RBS | | | | MATH | | | |
|------|--------------------|--------------------|--------------|--------------|--------------------|--------------------|--------------|--------------|--------------------|--------------------|--------------|--------------|
| | $\delta_{avg}(\%)$ | $\delta_{max}(\%)$ | $T_{avg}(s)$ | $T_{max}(s)$ | $\delta_{avg}(\%)$ | $\delta_{max}(\%)$ | $T_{avg}(s)$ | $T_{max}(s)$ | $\delta_{avg}(\%)$ | $\delta_{max}(\%)$ | $T_{avg}(s)$ | $T_{max}(s)$ |
| 120 | 0.167 | 1.759 | 0.02 | 1.00 | 0.010 | 0.091 | 7.57 | 11.00 | 0.000 | 0.000 | 21.17 | 96.00 |
| 140 | 0.162 | 1.585 | 0.01 | 1.00 | 0.010 | 0.110 | 13.09 | 17.00 | 0.000 | 0.088 | 29.90 | 122.00 |
| 160 | 0.171 | 1.421 | 0.02 | 1.00 | 0.010 | 0.066 | 20.38 | 26.00 | 0.000 | 0.000 | 38.97 | 205.00 |
| 180 | 0.132 | 0.955 | 0.03 | 1.00 | 0.009 | 0.065 | 31.25 | 38.00 | 0.000 | 0.000 | 49.85 | 151.00 |
| 200 | 0.138 | 1.172 | 0.08 | 1.00 | 0.008 | 0.055 | 55.70 | 90.00 | 0.000 | 0.000 | 92.54 | 296.00 |
| 300 | 0.101 | 1.185 | 0.25 | 1.00 | 0.005 | 0.032 | 195.52 | 238.00 | 0.001 | 0.158 | 214.79 | 301.00 |
| 500 | 0.046 | 1.090 | 1.53 | 9.00 | --- | --- | --- | --- | --- | --- | --- | --- |
| 1000 | 0.036 | 0.788 | 20.76 | 75.00 | --- | --- | --- | --- | --- | --- | --- | --- |
| 1500 | 0.022 | 0.754 | 94.89 | 496.00 | --- | --- | --- | --- | --- | --- | --- | --- |
| 2000 | 0.022 | 0.717 | 314.98 | 917.00 | --- | --- | --- | --- | --- | --- | --- | --- |
| 2500 | --- | --- | --- | --- | --- | --- | --- | --- | --- | --- | --- | --- |

Table 3: Comparison of the heuristics with best known solutions

More generally, our work is a proof of concept for the ML for OR paradigm of [12]: our results on $1|r_j|\sum_j C_j$ show that algorithms obtained using this ML for OR paradigm can outperform state-of-the-art algorithms on classic problems of the OR literature – the state-of-the-art algorithms for the applications considered in [12] were not as challenging. Furthermore, our results on $1|r_j|\sum_j C_j$ show the relevance of three extensions of the paradigm of [12] introduced in this paper. First, the Fenchel-Young loss approach we introduce for the learning problem leads to practically efficient algorithms. This paves the way to new applications of the paradigm because this learning method is generic. And, we can boost the performance of the algorithms obtained with the ML for OR paradigm by using local search heuristics as decoding algorithms as in `IMLH`, and by perturbing θ as in `itMLH`. These new ideas are not specific to the $1|r_j|\sum_j C_j$ problem, and future directions include their applications to other operations research problems.

Acknowledgments. We are grateful to Federico Della Croce for his advice on the algorithms to use in our benchmark, and to Fabio Salassa for providing us the code of the matheuristic.

References

- [1] A. Rinnooy Kan, Machine sequencing problem: classification, complexity and computation, Springer-Verlag, 1976.
- [2] K. Baker, Introduction to Sequencing and Scheduling, Wiley, 1974.
- [3] T. V.-D. F. Shang, L., Branch & memorize exact algorithms for sequencing problems: efficient embedding of memorization into search trees, *Computers & Operations Research* to appear (2020).
- [4] F. S. Tanaka, S., A dynamic-programming-based exact algorithm for general single-machine scheduling with machine idle time, *Journal of Scheduling* 15 (2012) 347–361.
- [5] T. R.-U. R. Chand, S., An iterative heuristic for the single machine dynamic total completion time scheduling problem, *Annals of Operations Research* 23 (1996) 641–651.
- [6] T. V. Della Croce, F., A recovering beam search algorithm for the one-machine dynamic total completion time scheduling problem, *Journal of the Operational Research Society* 53 (2002) 1275–1280.
- [7] V. T. Federico Della Croce, Fabio Salassa, A hybrid heuristic approach for single machine scheduling with release times, *Computers & Operations Research* 45 (2014) 7–11.
- [8] Y. Bengio, A. Lodi, A. Prouvost, Machine Learning for Combinatorial Optimization: A Methodological Tour d’Horizon, arXiv:1811.06128 [cs, stat] (2020). arXiv:1811.06128.

- [9] O. E. C. T. Asta, S., A tensor based hyper-heuristic for nurse rostering, Knowledge-Based Systems 98 (2016) 185–199.
- [10] M.-J. Y. N. A. Fonseca-Reyna, Y.C., Q-learning algorithm performance for m-machine, n –jobs flow shop scheduling problems to minimize makespan, Revista Investigacion Operacional 38 (2017) 281–290.
- [11] S. K. Arnold, F., What makes a solution good? the generation of problem-specific knowledge for heuristics, Computers & Operations Research 106 (2019) 280–288.
- [12] A. Parmentier, Learning to approximate industrial problems by operations research classic problems, 2020.
- [13] S. Nowozin, Structured Learning and Prediction in Computer Vision, Foundations and Trends® in Computer Graphics and Vision 6 (2010) 185–365. doi:10.1561/06000000033.
- [14] T.-Y. Liu, Learning to Rank for Information Retrieval, Springer, Berlin, 2011.
- [15] M. Blondel, A. F. T. Martins, V. Niculae, Learning with Fenchel-Young losses, Journal of Machine Learning Research 21 (2020) 1–69.
- [16] Q. Berthet, M. Blondel, O. Teboul, M. Cuturi, J.-P. Vert, F. Bach, Learning with Differentiable Perturbed Optimizers, arXiv:2002.08676 [cs, math, stat] (2020). arXiv:2002.08676.
- [17] A. S. Lewis, M. L. Overton, Nonsmooth optimization via quasi-Newton methods, Mathematical Programming 141 (2013) 135–163. doi:10.1007/s10107-012-0514-2.
- [18] P. Wieschollek, Cppoptimizationlibrary, <https://github.com/PatWie/CppNumericalSolvers>, 2016.

Appendix A. Proofs

Proof of Proposition 1. The main statement and the case where Z has positive and differentiable density $d\mu(z) \propto \exp(-\nu(z))dz$ summarize results in Sections 2 and 3 of Berthet et al. [16]. The sampled minimization case immediately follows from the fact that F^*, x as a maximum of m linear mappings. \square

Proposition 3 (Proof of Proposition 2). The convexity of L is an immediate corollary of the convexity of F established in Proposition 1. The Fenchel-Young inequality implies that the Fenchel-Young loss is non-negative, and equal to 0 only if $\phi(y; x)$ is in the subdifferential of F . Since, a sum of convex functions among which at least one is strictly convex, the first point follows from the results in Section 4 of [16][Section 4]. The second point immediately follows from the fact that $L(\theta; y, x) = 0$ if and only if $\phi(y; x) \in \partial_{\theta} F(\theta; x)$ and the fact that $\delta_{\theta} \max_{\eta \in \mathcal{C}(x)} \langle \theta + z | \eta \rangle$ is the set of η realizing the maximum.

Appendix B. Influence of perturbation strength

In this appendix, we explain why the strength of the perturbation does not impact the result of our method. We therefore suppose to use the perturbation $\boldsymbol{\theta} + \varepsilon \mathbf{Z}$ wherever we used the perturbation $\boldsymbol{\theta} + \mathbf{Z}$ in the paper, with $\varepsilon > 0$. Remark that this latter assumption is w.l.o.g.: Since $\mathbf{Z} \sim \mathcal{N}(0, I_d)$ is symmetric, moving from ε to $-\varepsilon$ does not change the distribution of $\boldsymbol{\theta} + \varepsilon \mathbf{Z}$. Let

$$F_\varepsilon(\boldsymbol{\theta}; x) := \mathbb{E}_{\mathbf{Z}} \left(\max_{y \in \mathcal{Y}(x)} \langle \boldsymbol{\theta} + \varepsilon \mathbf{Z} | \boldsymbol{\phi}(y; x) \rangle \right) =, \mathbb{E}_{\mathbf{Z}} \left(\max_{\boldsymbol{\eta} \in \mathcal{C}(x)} \langle \boldsymbol{\theta} + \varepsilon \mathbf{Z} | \boldsymbol{\eta} \rangle \right) \quad \text{and,}$$

$$\boldsymbol{\phi}_\varepsilon^*(\boldsymbol{\theta}; \mathbf{x}) := \mathbb{E}_{\mathbf{Z}} \left(\phi \left(\arg \max_{y \in \mathcal{Y}(x)} \langle \boldsymbol{\theta} + \varepsilon \mathbf{Z} | \boldsymbol{\phi}(y; x) \rangle; x \right) \right) = \mathbb{E} \left(\arg \max_{\boldsymbol{\eta} \in \mathcal{C}(x)} \langle \boldsymbol{\theta} + \varepsilon \mathbf{Z} | \boldsymbol{\eta} \rangle \right).$$

Let $F_\varepsilon^*(\cdot; x)$ be the Fenchel dual of $F_\varepsilon(\cdot; x)$ and recall that $\Omega := F_1^*$, i.e.,

$$F_\varepsilon^*(\boldsymbol{\eta}; x) := \arg \max_{\boldsymbol{\theta} \in \mathcal{C}(x)} \langle \boldsymbol{\eta} | \boldsymbol{\theta} \rangle - F_\varepsilon(\boldsymbol{\eta}; x) \quad \text{and} \quad \Omega(\boldsymbol{\eta}; x) := F_1^*(\boldsymbol{\eta}; x).$$

Proposition 2.2 of Berthet et al. [16] establishes that

$$F_\varepsilon(\boldsymbol{\theta}; x) = \varepsilon F_1\left(\frac{\boldsymbol{\theta}}{\varepsilon}\right), \quad F_\varepsilon^*(\cdot; x) = \varepsilon \Omega(\cdot; x), \quad \text{and} \quad \boldsymbol{\phi}_\varepsilon^*(\boldsymbol{\theta}; \mathbf{x}) = \boldsymbol{\phi}_1^*\left(\frac{\boldsymbol{\theta}}{\varepsilon}; \mathbf{x}\right). \quad (\text{B.1})$$

Replacing $\boldsymbol{\theta} + \mathbf{Z}$ by $\boldsymbol{\theta} + \varepsilon \mathbf{Z}$ amounts to replace F and Ω by F_ε and F_ε^* in the learning Problem (6). If we denote by $\boldsymbol{\theta}^*$ the optimal solution of Problem (6) with the standard perturbation $\varepsilon = 1$, and $\boldsymbol{\theta}_\varepsilon^*$ the optimal solution with the perturbation with strength ε , it follows from Equation (B.1) that $\boldsymbol{\theta}_\varepsilon^* = \varepsilon \boldsymbol{\theta}^*$. Hence, the structured prediction problem we obtain for any scenario ω is

$$\arg \max_{y \in \mathcal{Y}(x)} \langle \boldsymbol{\theta}_\varepsilon^* + \varepsilon \mathbf{Z}(\omega) | \boldsymbol{\phi}(y; x) \rangle = \arg \max_{y \in \mathcal{Y}(x)} \varepsilon \langle \boldsymbol{\theta}^* + \mathbf{Z}(\omega) | \boldsymbol{\phi}(y; x) \rangle$$

and we obtain the same predictions as the one we could have obtained using $\varepsilon = 1$.