A learning-based mathematical programming formulation for the automatic configuration of optimization solvers *

Gabriele Iommazzo^{1,2}, Claudia D'Ambrosio¹, Antonio Frangioni², and Leo Liberti¹

¹ LIX CNRS, École Polytechnique, Institut Polytechnique de Paris, Palaiseau, France, {dambrosio,giommazz,liberti}@lix.polytechnique.fr ² Dip. di Informatica, Università di Pisa, Pisa, Italy, frangio@di.unipi.it

Abstract. We propose a methodology, based on machine learning and optimization, for selecting a solver configuration for a given instance. First, we employ a set of solved instances and configurations in order to learn a performance function of the solver. Secondly, we formulate a mixed-integer nonlinear program where the objective/constraints explicitly encode the learnt information, and which we solve, upon the arrival of an unknown instance, to find the best solver configuration for that instance, based on the performance function. The main novelty of our approach lies in the fact that the configuration set search problem is formulated as a mathematical program, which allows us to a) enforce hard dependence and compatibility constraints on the configurations, and b) solve it efficiently with off-the-shelf optimization tools.

Keywords: automatic algorithm configuration, mathematical programming, machine learning, optimization solver configuration, hydro unit committment

1 Introduction

We address the problem of finding instance-wise optimal configurations for general Mathematical Programming (MP) solvers. We are particularly motivated by state-of-the-art general-purpose solvers, which combine a large set of diverse algorithmic components (relaxations, heuristics, cutting planes, branching, ...) and therefore have a long list of user-configurable parameters; tweaking them can have a significant impact on the quality of the obtained solution and/or on the efficiency of the solution process (see, e.g., [14]). Good solvers have effective default parameter configurations, carefully selected to provide good performances

^{*} This paper has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grant agreement n. 764759 "MINOA".

in most cases. Furthermore, tuning tools may be available (e.g., [18, Ch. 10]) which run the solver, with different configurations, on one or more instances within a given time limit, and record the best parameter values encountered. Despite all this, the produced parameter configurations may still be highly sub-optimal with specific instances. Hence, a manual search for the best parameter values may be required. This is a highly nontrivial and time-consuming task, due to the large amount of available parameters (see, e.g., [19]), which requires a profound knowledge of the application at hand and an extensive experience in solver usage. Therefore, it is of significant interest to develop general approaches, capable of performing it efficiently and effectively in an automatic way.

This setting is an instance of the Algorithm Configuration Problem (ACP) [27], which is defined as follows: given a target algorithm, its set of parameters, a set of instances of a problem class and a measure of the performance of the target algorithm on a pair (instance, algorithmic configuration), find the parameter configuration providing optimal algorithmic performance according to the given measure, on a specific instance or instance set. Several domains can benefit from automating this task. Some possible applications, beyond MP solvers (see, e.g., [12] and references therein), are: solver configuration for the propositional satisfiability problem, hyperparameter tuning of ML models or pipelines, algorithm selection, administering ad-hoc medical treatment, etc. Our approach for addressing the ACP on MP solvers is based on a two-fold process:

- (i) in the *Performance Map Learning Phase* (PMLP), supervised Machine Learning (ML) techniques [25] are used to learn a *performance function*, which maps some features of the instance being solved, and the parameter configuration, into some measure of solver efficiency and effectiveness;
- (ii) the formal model underlying the ML methodology used in the PMLP is translated into MP terms; the resulting formulation, together with constraints encoding the compatibility of the configuration parameter values, yields the *Configuration Set Search Problem* (CSSP), a Mixed-Integer Nonlinear Program (MINLP) which, for a given instance, finds the configuration providing optimal performance with respect to the performance function.

The main novelty of our approach lies in the fact that we explicitly model and optimize the CSSP using the mathematical description of the PMLP technique. This is in contrast to most of the existing ACP approaches, which instead employ heuristics such as local searches [2, 16], genetic algorithms [3], evolutionary strategies [8] and other methods [24]. Basically, most approaches consider the performance function as a black box, even when it is estimated by means of some ML technique and, therefore, they cannot reasonably hope to find a global minimum when the number of parameters grows. Rather, one of the strengths of our methodology is that it exploits the mathematical structure of the CSSP, solving it with sophisticated, off-the-shelf MP solvers. Moreover, formulating the CSSP by MP is advantageous as it allows the seamless integration of the compatibility constraints on the configuration parameters, which is something that other ACP methods may struggle with. The idea of using a ML predictor to define the unknown components (constraints, objective) of a MP has been already explored in data-driven optimization. In general, it is possible to represent the ML model of a mapping/relation as a MP (or, equivalently, a Constraint Programming model) and optimize upon this [23]; however, while this is in principle possible, the set of successful applications in practice is limited. Indeed, using this approach in the ACP context is, to the best of our knowledge, new; it also comes with some specific twists. We tested this idea with the following components: we configured nine parameters of the IBM ILOG CPLEX solver [18], which we employed to solve instances of the Hydro Unit Commitment (HUC) problem [7], we chose Support Vector Regression (SVR) [28] as the PMLP learning methodology, and we used the off-the-shelf MINLP solver Bonmin [5] to solve the CSSP.

The paper is structured as follows: in Sec. 2 we will review existing work on algorithm configuration; in Sec. 3 we will detail our approach and provide the explicit formulation of the CSSP with SVR; in Sec. 4 we will discuss some computational results.

2 The algorithm configuration problem

Most algorithms have a very high number of configurable parameters of various types (boolean, categorical, integer, continuous), which usually makes the ACP very hard to solve in practice. Notably, this issue significantly affects MP solvers: they are highly complex pieces of software, embedding several computational components that tackle the different phases of the solution process; the many available algorithmic choices are exposed to the user as a long list of tunable parameters (for example, more than 150 in CPLEX [19]).

Approaches to the ACP can be compared based on how they fit into the following two categories: Per-Set (PS) or Per-Instance (PI); offline or online. In PS approaches, the optimal configuration is defined as the one with the best overall performance over a set of instances belonging to the same problem class. Therefore, PS approaches first find the optimal configuration for a problem class and then use it for any instance pertaining to that class. The exploration of the configuration set is generally conducted by means of heuristics, such as various local search procedures [2, 16], racing methods [24], genetic algorithms [3] or other evolutionary algorithms [26]. In this context, an exception is, e.g., the approached described in [15], which predicts the performance of the target algorithm by random forest regression and then uses it to guide the sampling in an iterative local search. PS approaches, however, struggle when the target algorithm performance varies considerably among instances belonging to the same problem class. In these cases, PI methodologies, which assume that the optimal algorithmic configuration depends on the instance at hand, are likely to produce better configurations. However, while PS approaches are generally problem-agnostic, PI ones require prior knowledge of the problem at hand, to efficiently encode each instance by a set of features. PI approaches typically focus on learning a good surrogate map of the performance function, generally by performing regression: this approximation is used to direct the search in the configuration set. In [17], for example, linear basis function regression is used to

approximate the target algorithm runtime, which is defined as a map of both features and configurations; then, for a new instance with known features, the learnt map is evaluated at all configuration points in an exhaustive search, to find the estimated best one. However, other approaches may be used: in [8], for example, a map from instance features to optimal configuration is learnt by a neural network; in [6] the ACP is restricted to a single binary parameter of CPLEX, and a classifier is then trained to predict it. In [21], instead, CPLEX is run on a given instance for a certain amount of computational resources, then a ranking ML model is trained, on-the-fly, to learn the ordering of branch and bound variables, and it is then used to predict the best branching variable, at each node, for the rest of the execution. Another approach, presented in [20], first performs clustering on a set of instances, then uses the PS methodology described in [3] to find one good algorithmic configuration for each cluster. In [31], instead, instances are automatically clustered in the leaves of a trained decision tree, which also learns the best configuration for each leaf; at test time. a new instance is assigned to a leaf based on its features, and it receives the corresponding configuration. The purpose of an ACP approach is to provide a good algorithmic configuration upon the arrival of an unseen instance. We call a methodology offline if the learning happens before that moment, which is the case for all the approaches cited above. Otherwise, we call an ACP methodology online; these approaches normally use reinforcement learning techniques (see, e.g., [10]) or other heuristics [4].

In our approach, we define the performance of the target algorithm as a function of both features and controls, in order to account for the fact that the best configuration of a solver may vary among instances belonging to the same class of problems; this makes our approach PI. Moreover, we perform the PMLP only once, offline, which allows us to solve the resulting CSSP for any new instances in a matter of seconds. What makes our approach stand out from other methodologies is that the learning phase is treated as white-box: the prediction problem of the PMLP is formulated as a MP, which conveniently allows the explicit embedding of a mathematical encoding of the estimated performance into the CSSP, as its objective/constraints. This is opposed to treating the learned predictor as a black-box, and therefore using it as an oracle in brute-force searches or similar heuristics, that typically do not scale as well as optimization techniques.

3 The PMLP and the CSSP

Let \mathcal{A} be the target algorithm, and:

- $-C_{\mathcal{A}}$ be the set of feasible configurations of \mathcal{A} . We assume that each configuration $c \in C_{\mathcal{A}}$ can be encoded into a vector of binary and/or discrete values representing categorical and numerical parameters, and $C_{\mathcal{A}}$ can be described by means of linear constraints;
- Π be the problem to be solved, consisting of an infinite set of instances, and $\Pi' \subset \Pi$ be the (finite) set of instances used for the PMLP;

Learning-based MP formulation for the configuration of optimization solvers

- $-F_{II}$ be the set of feature vectors used to describe instances, encoded by vectors of continuous or discrete/categorical values (in the latter case they are labelled by reals);
- $-p_{\mathcal{A}}: F_{\Pi} \times \mathcal{C}_{\mathcal{A}} \longrightarrow \mathbb{R}$ be the performance function which maps a pair (f, c)(instance feature vector, configuration) to the outcome of an execution of \mathcal{A} (say in terms of the integrality gap reported by the solver after a time limit, but other measures are possible).

With the above definitions, the PMLP and the CSSP are detailed as follows.

3.1 Performance Map Learning Phase

In the PMLP we use a supervised ML predictor, e.g., SVR, to learn the coefficient vector θ^* providing the parameters of a prediction model $\bar{p}_{\mathcal{A}}(\cdot, \cdot, \theta) : F_{\Pi} \times \mathcal{C}_{\mathcal{A}} \to \mathbb{R}$ of the performance function $p_{\mathcal{A}}(\cdot, \cdot)$. The training set for the PMLP is

$$S = \left\{ (f_i, c_i, p_{\mathcal{A}}(f_i, c_i)) \mid i \in \{1 \dots s\} \right\} \subseteq F_{\Pi'} \times \mathcal{C}_{\mathcal{A}} \times \mathbb{R}, \tag{1}$$

where $s = |\mathcal{S}|$ and the training set labels $p_{\mathcal{A}}(f_i, c_i)$ are computed on the training vectors (f_i, c_i) . The vector θ^* is chosen as to hopefully provide a good estimate of $p_{\mathcal{A}}$ on points that do not belong to \mathcal{S} , with the details depending on the selected ML technology.

3.2 Configuration Space Search Problem

For a given instance f and parameter vector θ , $\mathsf{CSSP}(f,\theta)$ is the problem of finding the configuration with best estimated performance $\bar{p}_{\mathcal{A}}(f,c,\theta)$:

$$\mathsf{CSSP}(f,\theta) \equiv \min_{c \in \mathcal{C}_A} \bar{p}_{\mathcal{A}}(f,c,\theta) \,. \tag{2}$$

The actual implementation of $\mathsf{CSSP}(f,\theta)$ depends on the MP formulation selected to encode $\bar{p}_{\mathcal{A}}$, which may require auxiliary variables and constraints to define the properties of the ML predictor. If $\bar{p}_{\mathcal{A}}$ yields an accurate estimate of $p_{\mathcal{A}}$, we expect the optimum c^*_{cssp} of $\mathsf{CSSP}(f,\theta)$ to be a good approximation of the true optimal configuration c^* for solving f. However, we remark that a) $\mathsf{CSSP}(f,\theta)$ can be hard to solve, and b) it needs to be solved quickly (otherwise one might as well solve the instance f directly). Hence, incurring the additional computational overhead for solving the CSSP may be advantageous only when the instance at hand is "hard". Achieving a balance between PMLP accuracy and CSSP cost is one of the challenges of this research.

4 Experimental results

We tested our approach on 250 instances of the HUC problem and on 9 parameters of CPLEX, version 12.7. The PMLP and CSSP experiments were conducted on an Intel Xeon CPU E5-2620 v4 @ 2.10GHz architecture, while CPLEX was run on an Intel Xeon Gold 5118 CPU @ 2.30GHz. The pipeline was implemented in Python 3.6.8 [29] and AMPL Version 20200110 [13]. In the following, we detail the algorithmic set-up that we employed.

4.1 Building the dataset

- 1. Features. The HUC is the problem of finding the optimal scheduling of a pump-storage hydro power station, where the commitment and the power generation of the plant must be decided in a short term period, in which inflows and electricity prices are previously forecast. The goal is to maximize the revenue given by power selling (see, e.g., [1]). The time horizon is fixed to 24h and the underlying hydro system is also fixed, so that all the instances have the same size. Thus, only 54 elements that vary from day to day are features: the date, 24 hourly prices, 24 hourly inflows, initial and target water volumes, upper and lower bound admitted on the water volumes. We encode them in a vector f of 54 continuous/discrete components. All the instances have been randomly generated with an existing generator that accurately reproduces realistic settings.
- 2. Configuration parameters. Thanks to preliminary tests, we select a subset of 9 discrete CPLEX parameters (fpheur, dive, probe, heuristicfreq, startalgorithm and subalgorithm from mip.strategy; crossover from barrier; mircuts and flowcovers, from mip.cuts), for each of which we consider between 2 and 4 different values. We then combine them so as to obtain 2304 parameter configurations. A configuration is encoded by a vector $c \in \{0, 1\}^{23}$, where each categorical parameter is represented by its incidence vector.
- 3. Performance measure. We use the integrality gap to define $p_{\mathcal{A}}(f,c)$. It has been shown that MIP solvers can be affected by performance variability issues (see, e.g., [22]), due to executing the solver on different computing platforms, permuting rows/columns of a model, adding valid but redundant constraints, performing apparently neutral changes to the solution process, etc. In order to tackle this issue, first we sample three different random seeds. For each instance feature vector f and each configuration c, we then carry out the following procedure: (i) we run CPLEX (using the Python API) 3 times on the instance, using the different random seeds, for 60 seconds; (ii) we record the middle out of the three obtained performance values, to be assigned to the pair (f, c). At this point, our dataset contains $250 \times 2304 = 576000$ records. The performance measure thus obtained from CPLEX output, which we call $p_{cpx}(f,c)$, usually contains some extremely large floating point values (e.g., whenever the CPLEX gap has a value close to zero in the denominator), which unduly bias the learning process. We deal with this issue as follows: we compute the maximum \bar{p}_{cpx} , over all values of (the range of) $p_{\tt cpx},$ lower than a given threshold (set to 1e+5 in our experiments), re-set all values of p_{cpx} larger than the threshold to $\bar{p}_{cpx} + 100$, then rescale p_{cpx} so that it lies within the interval [0, 1]. The resulting performance measure, which in the following we call $p_{ml}(f,c)$, is also the chosen PMLP label. Moreover, we solve each HUC instance and we record the value of its optimum; then, $\forall (f, c)$, we compute the primal gap ρ_{prim} and the dual gap ρ_{dual} , i.e.: the distance between the optimal value of f and the value of the feasible solution found, and the distance between the value of the

tightest relaxation found and the optimal value (both over the optimum). We save p_{cpx} , p_{m1} , ρ_{prim} and ρ_{dual} in our dataset. We remark that setting the time-limit, imposed on CPLEX runs, to 60 seconds provides the solver enough time to move past the preliminary processing and to begin working on closing the gap, even for very hard instances (i.e., the ones with long preprocessing times); this allows us to measure the actual impact that different parameter configurations have on the chosen performance measure.

- 4. Feature engineering. We process the date in order to extract the season, the week-day, the year-day, two flags called isHoliday and isWeekend, and we perform several sine/cosine encodings, that are customarily used to treat cyclical features. Moreover, we craft new features by computing statistics on the remaining 54 features. This task takes around 12 minutes to complete for the whole data set.
- 5. Splitting the dataset. We randomly divide the instances into 187 In-Sample (IS) and 63 Out-of-Sample (OS), and split the dataset rows accordingly (430848 IS and 145152 OS). We use the IS data to perform Feature Selection (FS) and to train the SVR predictor; then, we assess the performance of the PMLP-CSSP pipeline both on OS instances, to test its generalization capabilities to unseen input, and on IS instances, to evaluate its performance on the data that we learn from, as detailed below.
- 6. Feature selection. We use Python's Pandas DataFrame's corr function to perform Pearson's Linear Correlation and sklearn RandomForestRegressor's feature_importances_ attribute to perform decision trees' Feature Importance, in order to get insights on which features contribute the most to yield accurate predictions. A detailed explanation of the employed FS techniques falls outside of the scope of this document. In the following, we use the shorthand "variables" to refer to the whole list of columns of the learning dataset. In order to perform FS, we use a dedicated subset of the IS dataset, composed of 19388 records and only employed for this task; performing the selected FS techniques on this dataset takes around 8 minutes, and reduces f to 22 components. For the configuration vectors we consider three FS scenarios: noFS, kindFS and aggFS, yielding c vectors with, respectively, 23, 14 and 10 components. We then filter the PMLP dataset according to the FS scenario at hand. However, after this filtering, the dataset may contain points with the same (f, c) but different labels $p_{m1}(f, c)$. Thus, for each instance: a) we delete the dataset columns that FS left out; b) we perform Pandas's group_by on the 22 columns chosen by FS, then c) compute the average p_{m1} of each group and use this as the new label (at this point, rows with the same (f, c) have the same label); d) we remove the duplicate rows of the dataset by Pandas drop_duplicates, keeping only one row. Lastly, we select ~ 11200 points for the PMLP.

4.2 PMLP experimental setup

The PMLP methodology of choice in this paper is SVR. Its advantages are: (a) the PMLP for training an SVR can be formulated as a convex Quadratic

Program (QP), which can be solved efficiently; (b) even complicated and possibly nonlinear performance functions can be learned by using the "kernel trick" [28]; (c) the solution of the PMLP for SVR provides a closed-form algebraic expression of the performance map \bar{p}_A , which yields an easier formulation of CSSP(f, θ). We use a Gaussian kernel during SVR training, which is the default choice in absence of any other meaningful prior [9]. We assess the prediction error of the predictor by Nested Cross Validation (NCV) [30]; furthermore, our training includes a phase for determining and saving the hyperparameters and the model coefficients of the SVR. These two tasks take approximately 4h in the aggFS and in the kindFS scenarios, and 5h in the noFS one. We use Python's sklearn.model_selection.RandomizedSearchCV for the inner loop of the NCV and a customized implementation for the outer loop, and sklearn.svm.SVR as the implementation of choice for the ML model.

A common issue in data-driven optimization is that using customary ML error metrics may not lead to good solutions of the optimization problem (see, for example, [11]). We tackled this issue by comparing the following metrics for the CV-based hyperparameter tuning phase, both computed on p_{m1} : the classical Mean Absolute Error $\mathsf{MAE} = \sum_{i \in S} |p_i - \bar{p}_i|$, where $p_i = p_{m1}(f_i, c_i)$ and $\bar{p}_i = \bar{p}_{\mathcal{A}}(f_i, c_i)$; the custom metric $\mathsf{cMAE}_{\delta} = \sum_{i \leq s} L_{\delta}(p_i, \bar{p}_i), \delta \in \{0.2, 0.3, 0.4\}$, where

$$L_{\delta}(p_{i},\bar{p}_{i}) = \begin{cases} (\bar{p}_{i}-p_{i}) \cdot (1+\frac{1}{1+exp(\bar{p}_{i}-\bar{p}_{i})}) & \text{if } p_{i} \leq \delta \text{ and } \bar{p}_{i} > p_{i} \\ (p_{i}-\bar{p}_{i}) \cdot (1+\frac{1}{1+exp(\bar{p}_{i}-p_{i})}) & \text{if } p_{i} \geq 1-\delta \text{ and } \bar{p}_{i} < p_{i} \\ (p_{i}-\bar{p}_{i}) & \text{if } \delta \leq p_{i} \leq 1-\delta \\ 0 & \text{otherwise.} \end{cases}$$

4.3 CSSP experimental setup

The choice of a Gaussian kernel in the SVR formulation makes the CSSP a MINLP with a nonconvex objective function $\bar{p}_{\mathcal{A}}$. More precisely, for a given instance with features \bar{f} , our CSSP is:

$$\min_{c \in \mathcal{C}_{\mathcal{A}}} \sum_{i=1}^{s} \alpha_i \exp\left(-\gamma \|(f_i, c_i) - (\bar{f}, c)\|_2^2\right)$$
(3)

where, for all $i \leq s$, (f_i, c_i) belong to the training set, α_i are the dual solutions of the SVR, γ is the scaling parameter of the Gaussian kernel, and C_A is defined by mixed-integer linear programming constraints encoding the dependences/compatibility of the configurations. We use AMPL to formulate the CSSP, and the nonlinear solver Bonmin [5], manually configured (with settings heuristic_dive_fractional yes, algorithm B-Hyb, heuristic_feasibility _pump yes) and with a time limit of 60 seconds, to solve it; then we retrieve, for each instance f, the Bonmin solution c_{bm}^* . Since we have enumerated all possible configurations, we can also compute the "true" global optimum $c_{cssp}^* = \arg\min\{\bar{p}_A(f,c), c \in C_A\}$ for sake of comparison. In Table 1, we report the percentage of cases where $c_{bm}^* = c_{cssp}^*$ ("%glob. mins") and, for all the instances where this is not true, the average distance between $\bar{p}_A(c_{bm}^*)$ and $\bar{p}_A(c_{cssp}^*)$, over

	set	FS	%glob. mins	avg loc. mins	CSSP time				
-		noFS	83.69	$2.013\mathrm{e}{-02}$	15.36				
	\mathbf{IS}	kindFS	87.70	2.758e - 02	10.73				
		aggFS	84.49	3.122e - 02	5.95				
-		noFS	85.32	$1.594\mathrm{e}{-02}$	13.44				
	OS	kindFS	90.48	2.014e - 02	12.23				
		aggFS	93.25	1.957e - 02	6.15				
Table 1. Quality of Bonmin's solutions, w.r.t. \bar{p}									

all the instances of the considered set ("avg loc. mins"); we also report the average CSSP solution time ("CSSP time"). The kindFS and the aggFS scenarios achieve better results, in terms of "%glob mins", than the noFS one. Furthermore, the local optima found by Bonmin are quite good ones: they are never larger than 3.2% (see "avg loc. mins") of c_{cssp}^* . The time that Bonmin takes to solve the CSSP is reduced from the noFS scenario to the aggFS one. This is due to the fact that the kindFS and the aggFS CSSP formulations have less variables than the noFS one, and so they are easier to solve. Bonmin needs, on average, less than 16 seconds to solve any CSSP; however, devising more efficient techniques to solve the CSSP (say, reformulations, decomposition, ...) might be necessary if our approach is scaled to considerably more algorithmic parameters.

4.4 Results

In order to assess the performance of the approach, we retrieve $p_{cpx}(c_{bm}^*)$, $p_{cpx}(c_{cpx})$ (CPLEX default configuration), the primal and dual gap of c_{bm}^* and c_{cpx} from the filtered dataset, for every IS and OS instance. Tab. 2 shows: the wins "%w" and

set	FS	%w	%w+d	%wnond	avg d	avg w	avg l	
	noFS	47.06	96.12	92.39	0	3.400e + 14	6.493e + 19	,
IS	kindFS	48.91	98.02	96.12	0	3.415e + 14	$8.895e{+}19$	
	aggFS	48.98	97.79	95.70	0	3.438e + 14	8.530e + 19	
	noFS	33.73	83.73	67.50	0	3.148e + 14	7.042e+19	-
OS	kindFS	36.64	82.61	68.06	0	$3.051e{+}14$	$5.385e{+19}$	
	aggFS	33.86	77.05	59.78	0	$3.335\mathrm{e}{+14}$	$4.486e{+19}$	1
D.	·			. /	* \	1 50		TO

Table 2. Pipeline quality w.r.t. $p_{cpx}(c_{bm}^*)$, by FS scenario and IS/OS set

the non-worsenings "%w+d", i.e., the percentage of instances such that $p_{cpx}(c_{bm}^*)$ is $< \text{or} \leq \text{than } p_{cpx}(c_{cpx})$, by the first sixteen decimal digits of p_{cpx} , in scientific notation; the wins-over-nondraws "%w_{nond}", i.e., the percent wins over the instances such that $p_{cpx}(c_{bm}^*) \neq p_{cpx}(c_{cpx})$; the average $|p_{cpx}(c_{cpx}) - p_{cpx}(c_{bm}^*)|$, over all the instances which score a win ("avg w") or a loss ("avg l"); the average $|p_{cpx}(c_{bm}^*) - \underline{p}_{cpx}|$ over all the other instances, where $\underline{p}_{cpx} = \min_{c \in C_A} p_{cpx}(f, c)$ for a given f ("avg d"). The "%w", "%w+d" and "%w_{nond}" are higher on IS instances than on the OS ones. The IS instances are used as the training set, so it is not surprising that \bar{p}_A is less accurate at OS instances; this results in worse CSSP solutions for OS instances. The fact that "avg d" is always 0 implies that, whenever CPLEX — configured by c_{cpx} — can close the gap, our c_{bm}^* proves to be just as efficient. In the nondraws, our approach shows consistent gains with respect to c_{cpx} , both on IS and OS instances. From this we gather that \bar{p}_A provides an accurate approximation of p_{cpx} 's global minima, even at points outside the training set. The noFS scenario presents the worst results. The aggFS

scenario achieves the highest percent wins on IS instances; it also provides the largest average wins and the smallest average loss, overall. However, the kindFS scenario presents the best performance, by the highest "%w" on OS instances and by the best overall "%w_{nond}". In Tab. 3, we report the percentage of IS

$ \begin{array}{ c c c c c c c c c c c c c c c c c c c$	set	\mathbf{FS}	metric	% feas c_{bm}^*	%feas c_{cpx}	ℓ ^{bm} prim	$\rho_{\rm prim}^{\rm cpx}$	ρ_{dual}^{bm}	ρ_{dual}^{cpx}
$IS \begin{array}{ c c c c c c c c c c c c c c c c c c c$			cMAE.2	96.79		1.504e - 01		9.708e - 02	
$IS \begin{array}{c c c c c c c c c c c c c c c c c c c $		noFS	cMAE.3	95.72	100	1.193e - 01	5.812e-01	8.104e - 02	1.018e-01
$\begin{tabular}{ c c c c c c c c c c c c c c c c c c c$			cMAE.4	98.40		1.143e - 01		7.253e - 02	
$\begin{tabular}{ c c c c c c c c c c c c c c c c c c c$			MAE	98.93		1.058e - 01		4.698e - 02	
$\begin{tabular}{ c c c c c c c c c c c c c c c c c c c$	IS		cMAE.2	98.75		6.975e - 02		3.792e - 02	
$ \begin{array}{c c c c c c c c c c c c c c c c c c c $			cMAE.3	99.38		8.271e - 02		4.245e - 02	
$\begin{tabular}{ c c c c c c c c c c c c c c c c c c c$		kindFS	cMAE.4	98.48		7.085e - 02		5.015e - 02	
$\begin{tabular}{ c c c c c c c c c c c c c c c c c c c$			MAE	96.43		7.183e - 02		4.159e - 02	
cMAE_3 97.86 7.669e-02 $5.111e-02$ aggFS cMAE_4 97.59 $9.267e-02$ $6.680e-02$ MAE 97.59 $7.747e-02$ 4.998e-02			cMAE.2	98.04		7.850e - 02		5.179e - 02	
aggFS cMAE.4 97.59 9.267e-02 6.680e-02 MAE 97.59 7.747e-02 4.998e-02			cMAE.3	97.86		7.669e - 02		5.111e - 02	
MAE 97.59 7.747e-02 4.998e-02		aggFS	cMAE.4	97.59		9.267e - 02		6.680e - 02	
			MAE	97.59		7.747e - 02		4.998e - 02	
cMAE.2 87.30 6.968e-02 7.931e-02			cMAE.2	87.30	100	6.968e - 02	4.633e-01	7.931e-02	8.121e-02
cMAE _{.3} 92.06 1.231e-01 8.500e-02		noFS	cMAE.3	92.06		1.231e - 01		8.500e - 02	
noFS cMAE.4 88.89 1.305e-01 9.687e-02			cMAE.4	88.89		1.305e - 01		9.687e - 02	
MAE 93.65 1.371e-01 6.611e-02			MAE	93.65		1.371e - 01		6.611e - 02	
cMAE.2 89.68 1.471e-01 6.471e-02			cMAE.2	89.68		1.471e - 01		6.471e - 02	
cMAE.3 87.83 1.261e-01 1.189e-01			cMAE.3	87.83		1.261e - 01		1.189e - 01	
OS kindFS cMAE _{.4} 91.01 100 1.541e-01 4.633e-01 1.283e-01 8.121e-0	OS	kindFS	cMAE.4	91.01		1.541e - 01		1.283e - 01	
MAE 92.86 1.486e-01 1.078e-01	00		MAE	92.86		1.486e - 01		1.078e - 01	
cMAE.2 89.42 2.054e-01 9.062e-02		aggFS	cMAE.2	89.42		2.054e - 01		9.062e - 02	
CMAE _{.3} 88.89 2.061e-01 9.395e-02			cMAE.3	88.89		2.061e - 01		9.395e - 02	
aggrS cMAE.4 88.10 2.075e-01 9.874e-02			cMAE.4	88.10		2.075e - 01		9.874e - 02	
MAE 88.62 1.792e-01 7.733e-02			MAE	88.62		1.792e - 01		7.733e - 02	

Table 3. Quality of the solutions attained by CPLEX, configured by c_{bm}^* and c_{cpx} , solving HUC instances, aggregated by IS/OS set, FS scenario and PMLP metric

and OS instances such that CPLEX, configured by c_{bm}^* and c_{cpx} , manages to find a feasible solution ("%feas c_{bm}^* " and "%feas c_{cpx} "); for those instances, the columns " ρ_{prim}^{bm} ", " ρ_{qrim}^{cpx} ", " ρ_{dual}^{bm} " and " ρ_{dual}^{cpx} " report the average primal and dual gap achieved by the solver. CPLEX's default configuration always allows the solver to obtain a feasible solution within the time-limit. Our approach presents similar results on IS instances ("%feas c_{bm}^* " is approximately 98%), but it has slightly worse performances on the OS ones ("%feas c_{bm}^* " is around 90%). However, the primal/dual gaps provided by our methodology are always better than those achieved by using CPLEX's default setting; actually, they are up to an order of magnitude smaller, in IS instances. The cMAE_{.4} and the MAE metrics provide the highest "%feas c_{bm}^* ", respectively, on IS instances (above 98%) and OS instances (around 92%). On average, the MAE also provides the best primal/dual gaps for IS instances, while the cMAE_{.4} is the best choice for OS instances. Lastly, while the kindFS scenario prevails on IS instances, the noFS one dominates on the OS ones.

5 Conclusions

The methodology presented in this paper conflates ML and MP techniques to solve the ACP. All in all, the results show that the approach is promising, in that the configurations that it provides typically have better primal/dual gaps than CPLEX's. The fact that, in a small fraction of the cases, no feasible solution is found depends on how infeasibility is encoded by p_{m1} ; an easy fix to this issue would be to use a performance measure promoting feasibility. Notably, it is interesting that, in some cases, the custom cMAE error metric outperforms the classical MAE. We also observed that using FS techniques is conducive to much easier CSSP formulations, without overly affecting the quality of the solutions. Overall, since choices taken at any point in the pipeline affect its final outcome, a number of details have to be carefully considered for the approach to work.

References

- van Ackooij, W., et al.: Shortest path problem variants for the hydro unit commitment problem. Electronic Notes in Discrete Mathematics 69, 309 – 316 (2018), joint EURO/ALIO International Conference 2018 on Applied Combinatorial Optimization (EURO/ALIO 2018)
- Adenso-Díaz, B., Laguna, M.: Fine-tuning of algorithms using Fractional Experimental Design and Local Search. Operations Research 54(1), 99–114 (2006)
- Ansótegui, C., Sellmann, M., Tierney, K.: A gender-based genetic algorithm for the automatic configuration of algorithms. In: Proceedings of the 15th International Conference on Principles and Practice of Constraint Programming. pp. 142–157. CP'09, Springer-Verlag, Berlin, Heidelberg (2009)
- 4. Battiti, R., Brunato, M.: Reactive Search: machine learning for memory-based heuristics. Tech. rep., University of Trento (2005)
- 5. Bonami, P., et al.: An algorithmic framework for convex mixed integer nonlinear programs. Discrete Optimization 5(2), 186–204 (2008)
- Bonami, P., Lodi, A., Zarpellon, G.: Learning a classification of mixed-integer quadratic programming problems. In: van Hoeve (eds), W. (ed.) Integration of Constraint Programming, Artificial Intelligence, and Operations Research. CPAIOR 2018. Lecture Notes in Control and Information Sciences, vol. vol. 10848, pp. 595– 604. Springer, Cham (2018)
- Borghetti, A., D'Ambrosio, C., Lodi, A., Martello, S.: An MILP approach for shortterm hydro scheduling and unit commitment with head-dependent reservoir. IEEE Transactions on Power Systems 23(3), 1115–1124 (2008)
- Brendel, M., Schoenauer, M.: Instance-based parameter tuning for Evolutionary AI Planning. In: Proceedings of the 13th Annual Conference Companion on Genetic and Evolutionary Computation. pp. 591–598. GECCO '11, ACM (2011)
- 9. Cristianini, N., Shawe-Taylor, J.: An Introduction to Support Vector Machines: And Other Kernel-based Learning Methods. Cambridge University Press (2000)
- Degroote, H., Bischl, B., Kotthoff, L., De Causmaecker, P.: Reinforcement Learning for automatic online algorithm selection - an empirical study. In: Proceedings of the 16th ITAT Conference Information Technologies - Applications and Theory. pp. 93–101 (2016)
- 11. Demirović, E., et al.: An investigation into prediction + optimisation for the knapsack problem. In: CPAIOR. pp. 241–257 (2019)
- 12. Eggensperger, K., Lindauer, M., Hutter, F.: Pitfalls and best practices in algorithm configuration. CoRR abs/1705.06058 (2017)
- 13. Fourer, R., Gay, D.: The AMPL Book. Duxbury Press, Pacific Grove (2002)
- 14. Hutter, F., Hoos, H.H., Leyton-Brown, K.: Automated configuration of mixed integer programming solvers. In: Proceedings of the 7th International Conference on

Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems. pp. 186–202. CPAIOR'10, Springer-Verlag, Berlin, Heidelberg (2010)

- Hutter, F., Hoos, H.H., Leyton-Brown, K.: Sequential Model-based Optimization for general algorithm configuration. In: Proceedings of the 5th International Conference on Learning and Intelligent Optimization. pp. 507–523. LION'05, Springer-Verlag (2011)
- Hutter, F., Hoos, H.H., Leyton-Brown, K., Stützle, T.: ParamILS: An automatic algorithm configuration framework. Journal of Artificial Intelligence Research 36(1), 267–306 (2009)
- Hutter, F., Youssef, H.: Parameter adjustment based on performance prediction: Towards an instance-aware problem solver. Tech. rep., In: Technical Report: MSR-TR-2005125, Microsoft Research (2005)
- IBM: IBM ILOG CPLEX Optimization Studio, CPLEX 12.7 User's Manual. IBM (2016)
- IBM: IBM ILOG CPLEX Optimization Studio CPLEX Parameters Reference (2016)
- Kadioglu, S., Malitsky, Y., Sellmann, M., Tierney, K.: ISAC: Instance Specific Algorithm Configuration. In: Proceedings of the 2010 Conference on ECAI 2010: 19th European Conference on Artificial Intelligence. pp. 751–756. IOS Press, Amsterdam, The Netherlands (2010)
- Khalil, E.B., Bodic, P.L., Song, L., Nemhauser, G., Dilkina, B.: Learning to branch in mixed integer programming. In: Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence. pp. 724–731. AAAI'16, AAAI Press (2016)
- Lodi, A., Tramontani, A.: Performance variability in mixed-integer programming. Tutorials in Operations Research, Vol. 10 pp. 1–12 (2013)
- Lombardi, M., Milano, M., Bartolini, A.: Empirical decision model learning. Artificial Intelligence 244, 343–367 (2017)
- López-Ibáñez, M., Dubois-Lacoste, J., Pérez Cáceres, L., Birattari, M., Stützle, T.: The irace package: iterated racing for automatic algorithm configuration. Operations Research Perspectives 3, 43–58 (2016)
- Mohri, M., Rostamizadeh, A., Talwalkar, A.: Foundations of Machine Learning. The MIT Press, 2nd edn. (2018)
- Nannen, V., Eiben, A.E.: Relevance estimation and value calibration of evolutionary algorithm parameters. In: Proceedings of the 20th International Joint Conference on Artifical Intelligence. pp. 975–980. IJCAI'07, Morgan Kaufmann Publishers Inc. (2007)
- Rice, J.R.: The algorithm selection problem. Advances in Computers 15, 65–118 (1976)
- Smola, A.J., Schölkopf, B.: A tutorial on support vector regression. Statistics and Computing 14(3), 199–222 (2004)
- Van Rossum, G., Drake, F.L.: Python 3 Reference Manual. CreateSpace, Scotts Valley, CA (2009)
- Varma, S., Simon, R.: Bias in error estimation when using cross-validation for model selection. BMC Bioinformatics 7(91) (2006)
- Vilas Boas, M.G., Gambini Santos, H., de Campos Merschmann, L.H., Vanden Berghe, G.: Optimal decision trees for the algorithm selection problem: Integer programming based approaches. CoRR abs/1907.02211 (2019)