



Warm-starting constraint generation for mixed-integer optimization: A Machine Learning approach



Asunción Jiménez-Cordero^{a,b,*}, Juan Miguel Morales^a, Salvador Pineda^a

^a OASYS Group, University of Málaga, Ada Byron Research Building, Arquitecto Francisco Peñalosa St., 18, 29010, Málaga, Spain

^b Department of Mathematical Analysis, Statistics and Operations Research and Applied Mathematics (Area of Statistics and Operations Research), University of Málaga, Ada Byron Research Building, Arquitecto Francisco Peñalosa St., 18, 29010, Málaga, Spain

ARTICLE INFO

Article history:

Received 12 May 2022

Received in revised form 14 July 2022

Accepted 27 July 2022

Available online 1 August 2022

Keywords:

Mixed integer linear programming

Machine learning

Constraint generation

Warm-start

Feasibility and optimality guarantees

ABSTRACT

Mixed Integer Linear Programs (MILP) are well known to be NP-hard (Non-deterministic Polynomial-time hard) problems in general. Even though pure optimization-based methods, such as constraint generation, are guaranteed to provide an optimal solution if enough time is given, their use in online applications remains a great challenge due to their usual excessive time requirements. To alleviate their computational burden, some machine learning techniques (ML) have been proposed in the literature, using the information provided by previously solved MILP instances. Unfortunately, these techniques report a non-negligible percentage of infeasible or suboptimal instances.

By linking mathematical optimization and machine learning, this paper proposes a novel approach that speeds up the traditional constraint generation method, preserving feasibility and optimality guarantees. In particular, we first identify offline the so-called invariant constraint set of past MILP instances. We then train (also offline) a machine learning method to learn an invariant constraint set as a function of the problem parameters of each instance. Next, we predict online an invariant constraint set of the new unseen MILP application and use it to initialize the constraint generation method. This warm-started strategy significantly reduces the number of iterations to reach optimality, and therefore, the computational burden to solve online each MILP problem is significantly reduced. Very importantly, all the feasibility and optimality theoretical guarantees of the traditional constraint generation method are inherited by our proposed methodology. The computational performance of the proposed approach is quantified through synthetic and real-life MILP applications.

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

Recent papers, e.g., [1–4], have shown the potential in combining Mathematical Optimization and Machine Learning. Particularly, Mixed Integer Linear Programming (MILP) is known to be a powerful and flexible tool for modeling and solving a wide variety of decision-making problems, as can be confirmed from [5,6]. However, most MILPs are known to be NP-hard (Non-deterministic Polynomial-time hard) and therefore, using them for online applications becomes challenging. For this reason, the design of novel Machine-Learning-assisted (ML) techniques that reduce the computational burden of MILPs has recently become a popular research topic. The works in [7,8] are a valid proof of that. Although different learning strategies can be considered, all methods assume that a set of slight variations of the same problem have been previously solved, and their input data and

solutions are available. This is a reasonable assumption since many optimization problems are frequently solved for a range of input parameters in online applications.

Several works have proposed approaches that seek to preserve optimality guarantees while substantially reducing the solution time. For instance, the authors of [9] design a data-driven methodology to improve the use of heuristics in branch-and-bound. To be more precise, the authors train a machine learning model that provides a schedule of heuristics, specifying when and for how long each heuristic is executed. They numerically show that this smart schedule is very likely to substantially diminish the time to solve the mixed-integer program to optimality. Even though the proposed method yields impressive results, collecting data from the branch-and-bound heuristics may be a difficult and solver-dependent task. Hence, alternative approaches that only need and learn from the information provided by the MILP optimal solution are also considered. One of the most commonly used strategies in this line consists in building, from this information, a *simpler* formulation of the original MILP that is faster to solve.

A pure optimization-based method that can be applied to this end is *constraint generation*. A detailed explanation of this

* Corresponding author.

E-mail addresses: asuncionjc@uma.es (A. Jiménez-Cordero), juanmi82mg@gmail.com (J.M. Morales), spinedamorente@gmail.com (S. Pineda).

methodology can be found in [10]. Essentially, it sequentially adds violated constraints until the optimal solution is found, at the expense of solving a possibly large number of MILPs. As a consequence, the computational burden associated with this strategy may be unacceptable for online applications. One way of alleviating such computational effort would be to provide a good warm-start. That is, if a good initial set of constraints is given, then just a few iterations of the constraint generation method are executed, reducing the computational time. In this vein, the authors of [11] propose a learning strategy that uses a modified nearest neighbor methodology to screen out superfluous constraints of the Unit Commitment (UC) problem. The reader is referred to [12] for further details about this problem. Similarly, the method presented in [13] selects a subset of constraints for the same problem. All these works execute the constraint generation algorithm with an initial set of constraints that is inferred only from the constraints that appeared binding in previous instances of the problem, that is, from constraints that held with equality at the optimal solution. Using these binding constraints is not enough to recover the optimal solution in MILPs, as shown in [11]. Indeed, when integer variables are involved in the optimization problem, the optimal solution must satisfy the constraints included in the so-called *invariant constraint set* defined in [14]. Apart from the constraints that hold with equality, such an invariant constraint set includes other non-binding constraints that are crucial to attaining the optimal solution.

Therefore, the contributions and objectives achieved in this paper are:

- The authors in [11] completely ignore the critical non-binding constraints. Hence, the learned warm-start set is not good enough and, to attain optimality, the constraint generation method may require a possibly large number of iterations, thus resulting in quite modest computational savings. In contrast, we show that effectively warm-starting the constraint generation method is essential to reaching the optimal solution of a MILP after just a few iterations of the algorithm.
- To this aim, we identify offline an invariant constraint set for each past instance. Then, we train a machine learning model of our choice that returns the prediction of an invariant constraint set from MILP parameters.
- Such a learned set is used to warm-start the constraint generation method of a new MILP. This way, the optimal solution is attained after running just a small number of iterations in the constraint generation method, and the computational time of the online MILP is considerably reduced, as empirically shown using synthetic and real-life instances”.

The remainder of this paper is structured as follows: Section 2 presents MILP notation and discusses the difficulties of computing an invariant constraint set when integer variables appear. Section 3 details the proposed methodology. Section 4 focuses on the numerical experiments. Finally, Section 5 provides the conclusions and future research lines.

2. Invariant constraint set in MILPs

A general MILP can be formulated as follows:

$$\begin{cases} \min_{\mathbf{z} \in \mathbb{R}^n \times \mathbb{Z}^q} & \mathbf{c}^\top \mathbf{z} \\ \text{s.t.} & \mathbf{a}_j^\top \mathbf{z} \leq b_j, \quad \forall j \in \mathcal{J} \end{cases} \quad (P_\theta[\mathcal{J}])$$

where $\mathbf{c}, \mathbf{a}_j \in \mathbb{R}^{n+q}, \forall j, b_j \in \mathbb{R}, \forall j$, are input parameters and $\mathbf{z} = (\mathbf{x}, \mathbf{y})$ is the decision variable vector formed by the continuous variables $\mathbf{x} \in \mathbb{R}^n$ and the integer variables $\mathbf{y} \in \mathbb{Z}^q$. For convenience, we collect all those input parameters into the set

θ , that is, $\theta = \{\mathbf{c}, \mathbf{a}_j, b_j, \forall j \in \mathcal{J}\}$, so that $(P_\theta[\mathcal{J}])$ denotes the optimization problem with the set of input parameters θ and the set of constraints \mathcal{J} . Note that there is some abuse of notation in $(P_\theta[\mathcal{J}])$. As it is written, \mathcal{J} and θ can change independently. However, this is not true, since the definition of θ explicitly depends on \mathcal{J} . Hence, to be rigorous, we should write $(P_{\theta_{\mathcal{J}}}[\mathcal{J}])$. Nevertheless, in order to make the notation clearer, we remove the subindex \mathcal{J} in $\theta_{\mathcal{J}}$. In addition, for simplicity, we assume that problem $(P_\theta[\mathcal{J}])$ is bounded and feasible, and that its optimal solution $\mathbf{z}_\theta^*[\mathcal{J}]$ is assumed to be unique. Note, however, that if multiple optimal solutions appear, retaining just one of them is enough for our proposal.

The feasible region defined by constraints in \mathcal{J} includes the subset of the so-called *binding constraints* \mathcal{B} . Particularly, \mathcal{B} is comprised of the inequality constraints that hold with equality at the optimal solution, i.e., $\mathcal{B} = \{j \in \mathcal{J} : \mathbf{a}_j^\top \mathbf{z}_\theta^*[\mathcal{J}] = b_j\}$. Besides, according to [14], a subset of constraints $\mathcal{S} \subset \mathcal{J}$ is defined to be an *invariant constraint set*, if the objective values of problems $(P_\theta[\mathcal{J}])$ and $(P_\theta[\mathcal{S}])$ coincide, i.e., if $\mathbf{c}^\top \mathbf{z}_\theta^*[\mathcal{J}] = \mathbf{c}^\top \mathbf{z}_\theta^*[\mathcal{S}]$. Note that, following the previous definition, a unique invariant constraint set may not exist. The relationship between these two sets of constraints, \mathcal{B} and \mathcal{S} , depends on whether problem $(P_\theta[\mathcal{J}])$ includes integer variables or not. Therefore, we discuss first the case in which all variables are continuous and then the more general case with both continuous and integer variables.

Let us first assume a particular case of $(P_\theta[\mathcal{J}])$ that only includes continuous variables, i.e., $q = 0$ and $\mathbf{z} = \mathbf{x} \in \mathbb{R}^n$. Using the optimal solution, $\mathbf{z}_\theta^*[\mathcal{J}]$, it is straightforward to determine the set of binding constraints \mathcal{B} . Actually, the invariant constraint sets \mathcal{S} of linear programming problems must contain all the constraints in \mathcal{B} , as the authors of [15] affirm. In other words, when the decision variables of the optimization problem are continuous, one can choose $\mathcal{S} = \mathcal{B}$. This way, instead of solving the original optimization problem $(P_\theta[\mathcal{J}])$, which may involve a high computational cost, the optimal solution is computed through the reduced problem $(P_\theta[\mathcal{S}]) = (P_\theta[\mathcal{B}])$. Such a reduced problem includes fewer constraints, typically implies a lower computational effort and is thus more appropriate for online applications.

Now take the more general case in which problem $(P_\theta[\mathcal{J}])$ includes both continuous and integer variables. In this case, authors of [11] demonstrate through an illustrative example that an invariant constraint set \mathcal{S} can also include constraints that are non-binding at the optimum. These constraints play a critical role when solving MILPs since they cannot be removed from the original feasible region without impairing the feasibility, and thus the optimality, of the so-obtained solution. In other words, when integer variables are involved in the optimization problem, an invariant constraint set includes not only binding constraints but also some of the non-binding ones, i.e., it holds that $\mathcal{S} \supset \mathcal{B}$. Like in the previous case, assume we have access to the optimal solution $\mathbf{z}_\theta^*[\mathcal{J}]$. While the set of binding constraints can also be easily determined by evaluating all constraints at the optimum, identifying such a subset of critical non-binding constraints is a more challenging task when integer variables appear. In addition, ignoring these non-binding constraints is quite dangerous since the optimal solution of the reduced problem may violate some of the constraints in the set \mathcal{J} , increasing the number of iterations required by the constraint generation method.

Our claim is that, if tuples $(\theta_t, \mathcal{S}_t)$ for previously solved instances $t = 1, \dots, T$ are used to train any machine learning algorithm $ML(\cdot)$, then the solution of the reduced formulation of a new unseen instance \tilde{t} with the predicted invariant constraint set, $P_{\theta_{\tilde{t}}}[\mathcal{S}_{\tilde{t}}]$ would more likely be feasible (and optimal) for the original problem $P_{\theta_{\tilde{t}}}[\mathcal{J}_{\tilde{t}}]$, than if only the tuples with the binding constraints $(\theta_t, \mathcal{B}_t), \forall t$ are considered, as done in [11]. As a consequence, the number of iterations to be run online in the constraint

generation method is decreased. Hence, it is desirable to warm-start the constraint generation algorithm with a constraint set as close to an invariant constraint set of the original MILP as possible. This way, the number of iterations executed by the algorithm is reduced, and so is its running time.

3. Methodology

The goal of this paper is to develop a data-driven approach that guarantees the optimal solution of an MILP in a reduced computational time. To this aim, we propose a methodology that efficiently warm-start the constraint generation method so that the number of iterations executed (online) is as small as possible. The key point of our strategy is to initialize (warm-start) the constraint generation method using a predicted invariant constraint set, $S \subset \mathcal{J}$, learned from past instances.

In particular, we propose a data-driven strategy that takes advantage of the information (θ_t, S_t) provided from the previously solved instances $t = 1, \dots, T$ to predict an invariant constraint set, $S_{\tilde{t}} = \text{ML}(\theta_{\tilde{t}})$ for a new unseen problem \tilde{t} , using a machine learning model of our choice, $\text{ML}(\cdot)$. See [16] for more details about the main machine learning tools. Then, once such a set $S_{\tilde{t}}$ is predicted, we only need to run a smaller number of iterations than those required for the original constraint generation method to converge to the optimal solution of the MILP. Indeed, if the predicted invariant constraint set exactly coincides with the actual one, then, only one iteration of the constraint generation method needs to be executed. In other words, when the prediction is perfect, then it suffices to solve $P_{\theta_{\tilde{t}}}[S_{\tilde{t}}]$.

As mentioned in Section 2, given an optimization problem, finding an invariant constraint set is a challenging task when integer variables appear. Developing an efficient procedure to learn an invariant constraint set in MILPs, $S \subset \mathcal{J}$, is a relevant research question that has not yet been properly answered in the literature. We propose in this paper a methodology that aims at determining, for each training instance t , an invariant constraint set S_t . The proposed approach to construct S_t for each instance t is also based on a constraint generation procedure. To compute an invariant constraint set S_t , for a previously solved instance t , we proceed as follows: we initialize the invariant constraint set S_t with the set of binding constraints \mathcal{B}_t . At each iteration, a reduced problem $P_{\theta_t}[S_t]$ is solved giving the optimal solution $\mathbf{z}_{\theta_t}^*[S_t]$. If all original constraints are satisfied, i.e., $\mathbf{a}_j^T \mathbf{z}_{\theta_t}^*[S_t] \leq b_j, \forall j \in \mathcal{J} \setminus S_t$, then the algorithm terminates. Otherwise, the most violated constraint is included in the set S_t , and a new iteration is run. This algorithm is run offline for each training instance t . Hence, adding, at each iteration, only the most violated constraint is appropriate for our proposal since the online running time is not affected. However, alternative strategies, such as including in the set S_t all the violated constraints at each iteration, can be considered. The pseudocode of the proposed procedure is given in Algorithm 1.

Algorithm 1 Identifying an invariant constraint set for each instance t

- 0) Initialize $S_t = \mathcal{B}_t$.
 - 1) Solve $P_{\theta_t}[S_t]$ with solution $\mathbf{z}_{\theta_t}^*[S_t]$.
 - 2) If $\max_{j \in \mathcal{J} \setminus S_t} \{\mathbf{a}_j^T \mathbf{z}_{\theta_t}^*[S_t] - b_j\} > 0$, go to step 3. Otherwise, stop.
 - 3) $S_t := S_t \cup \left\{ \arg \max_{j \in \mathcal{J} \setminus S_t} \{\mathbf{a}_j^T \mathbf{z}_{\theta_t}^*[S_t] - b_j\} \right\}$, go to step 1.
-

After running Algorithm 1, the information (θ_t, S_t) is available for all the instances $t = 1, \dots, T$ to train a machine learning model, $\text{ML}(\cdot)$, of our choice. In the next step, we take

the parameters $\theta_{\tilde{t}}$ of a new unseen problem instance \tilde{t} , and predict an invariant constraint set $S_{\tilde{t}} = \text{ML}(\theta_{\tilde{t}})$ with the already trained model. Finally, we run just a few iterations of the constraint generation method warm-started with the learned invariant constraint set $S_{\tilde{t}}$. The prediction of the invariant constraint set and the constraint generation method are executed online. In contrast, the strategy to build S_t for all training instances t (Algorithm 1), as well as the training of the machine learning algorithm $\text{ML}(\cdot)$ are performed offline. This way, the online computational burden is not affected. Algorithm 2 shows a pseudocode of the main steps of our approach.

Algorithm 2 Pseudocode of the proposed methodology.

Offline phase:

Input: $\{(\theta_t, \mathcal{B}_t)\}, \forall t$.

- 1) For each train instance t :
 - (a) Run Algorithm 1.
 - (b) Obtain an invariant constraint set, S_t .
- 2) Train $\text{ML}(\cdot)$ using $\{(\theta_t, S_t)\}, \forall t$.

Output: Trained ML methodology.

Online phase:

Input: Trained ML strategy of previous step and $\theta_{\tilde{t}}$ from a test instance \tilde{t} .

- 3) Predict $S_{\tilde{t}} = \text{ML}(\theta_{\tilde{t}})$.
- 4) Run CG initialized with the set $S_{\tilde{t}}$.

Output: Optimal solution of the test problem instance \tilde{t} .

The main advantages of the proposed methodology are described below:

- Since we are running a constraint generation procedure that is warm-started with a carefully built constraint set, our approach retains the convergence optimality guarantees from the standard constraint generation method.
- The procedure to build the invariant constraint sets, $S_t, \forall t$, of previously solved instances is also based on constraint generation. Hence, it is guaranteed to include all the non-binding constraints necessary to recover the optimal objective value. Therefore, all the past instances verify that the optimal solutions of the reduced problems are feasible and optimal for the original formulations.
- There is no condition about the machine learning algorithm that we apply in our approach. In other words, the sets of constraints S_t can be used for training *any* machine learning method.
- The invariant constraint sets, $S_t, \forall t$ and the machine learning algorithm, $\text{ML}(\cdot)$ are run offline. Therefore, the computational cost executed online to determine the optimal solution of a new unseen MILP instance is not affected.

4. Computational experiments

This section is devoted to the numerical experiments carried out in this paper. Section 4.1 details the experimental setup, whereas Section 4.2 explains the results derived from testing our proposal on two case studies.

4.1. Experimental setup

To show the efficiency of our approach, we compare it with two algorithms. The first one is the standard constraint generation algorithm, denoted as CG, which is based on pure optimization grounds and completely ignores the information provided

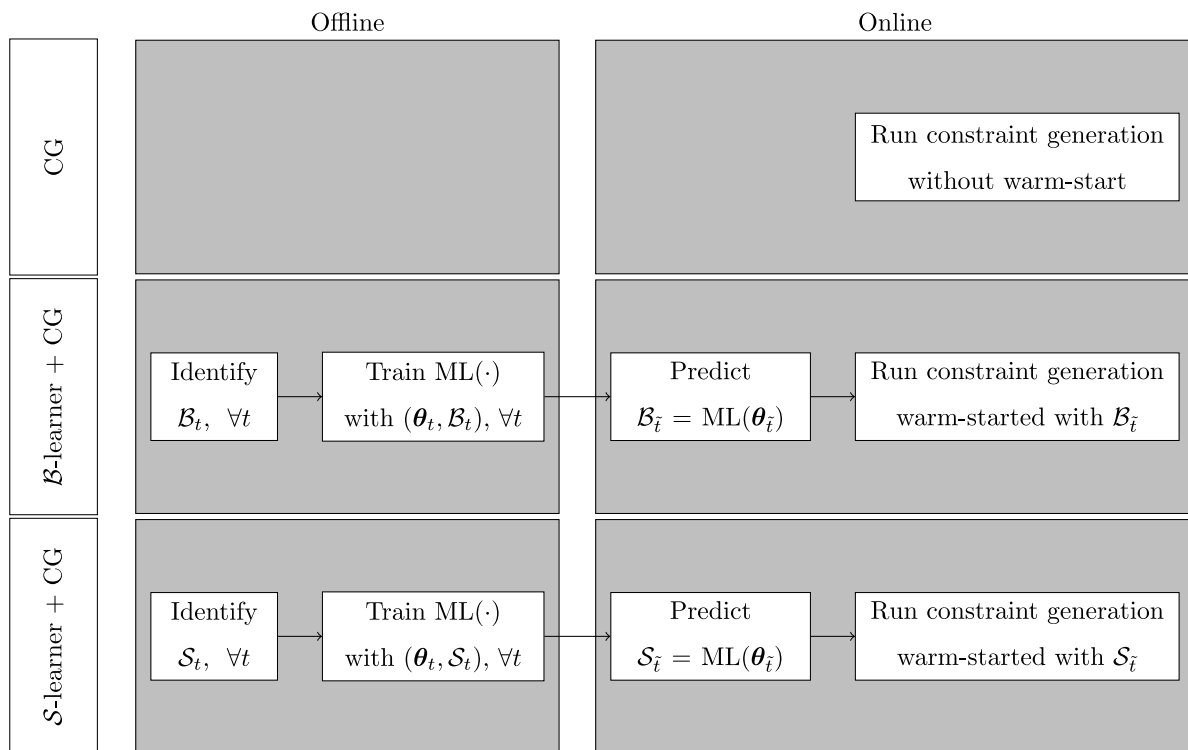


Fig. 1. Flowchart of the three algorithms.

by the data. In particular, for each instance, we sequentially add the violated constraints at each iteration. The second comparative approach is based on reference [11]. The authors propose a data-driven method where the constraint generation method is warm-started using only the information given by the binding constraint set, \mathcal{B} . We denote this method by \mathcal{B} -learner + CG. Finally, since we warm-start the constraint generation using an invariant constraint set \mathcal{S} , our methodology is denoted as \mathcal{S} -learner + CG.

It is well-known that machine learning performance highly depends on the data division into training and test samples. Thus, to get stable out-of-sample results, leave-one-out is executed in this paper. More details about this technique can be found in [17]. In particular, we assume given a database of T previously solved MILP instances. As mentioned in Section 2, we assume that the optimal solution of such instances is unique. Otherwise, retaining one of them is enough to our purposes. Note that, in the case of multiple solutions, solvers usually retain just one of them instead of the complete set of solutions. Hence, when multiple solutions appear, we only collect the one provided by the solver.

The leave-one-out strategy consists in running T iterations of our approach. At each iteration, we select one MILP instance and consider it as the test set, $\{\tilde{t}\}$ to be run in the online phase. The remaining $T - 1$ instances constitute the training set, $\{1, \dots, T\} \setminus \{\tilde{t}\}$. Then, we run the \mathcal{S} -learner + CG method. We first identify offline $T - 1$ invariant constraint sets \mathcal{S}_t , for all training instance t by running Algorithm 1, and train (also offline) a machine learning model, $\text{ML}(\cdot)$, that learns an invariant constraint set in terms of the MILP parameters. The next step (to be also performed online) is to use the already trained model to predict an invariant constraint set for the test instance. Finally, such a predicted set is utilized to warm-start the constraint generation procedure that results in the optimal solution of the test instance. For the sake of comparison, both data-based strategies are applied in an equivalent way. Naturally, the \mathcal{B} -learner + CG procedure is trained with the series of $T - 1$ binding constraints sets, \mathcal{B}_t ,

instead of the invariant constraint sets, \mathcal{S}_t . On the other hand, the pure optimization-based strategy CG does not need to divide the whole database into train and test sets. CG is run T times, one per MILP instance. Each time, an iterative algorithm that sequentially adds the violated constraints is run. Fig. 1 shows a scheme of our proposal (given in Algorithm 2) together with the two comparative algorithms, emphasizing the offline and online steps of each method.

The training of the machine learning model $\text{ML}(\cdot)$ and the subsequent construction of the set $\mathcal{S}_{\tilde{t}}$ (resp. $\mathcal{B}_{\tilde{t}}$) is addressed as a binary classification problem. For this purpose, we assign a label $s_t^j = 1$ or $s_t^j = -1$ to each constraint $j \in \mathcal{J}$ of each problem instance t with optimal solution $\mathbf{z}_{\theta_t}^*[\mathcal{J}]$, depending on whether that constraint is in \mathcal{S}_t (resp. \mathcal{B}_t) or not, respectively. Likewise, each constraint $j \in \mathcal{J}$ in the new problem instance \tilde{t} will be assigned the label $s_{\tilde{t}}^j = 1$ if the machine learning model $\text{ML}(\cdot)$ predicts that the constraint j is in $\mathcal{S}_{\tilde{t}}$ (resp. $\mathcal{B}_{\tilde{t}}$), and label $s_{\tilde{t}}^j = -1$ otherwise. Accordingly, any classification algorithm can, in principle, serve this purpose. In this paper, we select the well-known k nearest neighbors, knn, due to its simplicity. See Ref. [18] for more details in this regard. Nevertheless, alternative learning approaches such as Support Vector Machines, Neural Networks or Decision Trees can be applied as well. The work [17] explains the main properties of these methodologies.

Note that misclassifying has different consequences depending on the type of constraint we wrongly label. Indeed, adding a superfluous constraint into the set $\mathcal{S}_{\tilde{t}}$ is far much less damaging than failing to include a constraint in that set. The former case only leads to a slight increase in the size of the reduced problem, while the latter increases the risk that the solution to the reduced problem is infeasible in the original formulation, thus potentially increasing the number of iterations to be executed in the constraint generation method. For this reason, in this paper, we want to be on the conservative side when choosing the knn voting strategy. For a fixed k , we consider that a constraint j of an unseen test instance, \tilde{t} , belongs to an invariant constraint

set \mathcal{S}_i , i.e., $s_i^j = 1$, if and only if at least one of the k closest training instances includes such a constraint j in the set \mathcal{S}_i . To ensure a fair comparison, we apply the same voting strategy when constructing \mathcal{B}_i from \mathcal{B}_t , $\forall t$ using the knn classification. In any case, a feature of the knn method is that the larger the k , the larger the size of \mathcal{S}_i and \mathcal{B}_i . Thus, high values of k are expected to result in predicted sets \mathcal{S}_i and \mathcal{B}_i with a larger number of constraints. Therefore, the number of iterations in the constraint generation method may be reduced at the expense of a potential increase in the running time of each iteration. Then, there exists a trade-off between the value of k and the possible increment of the computational burden. The user should decide the value of k depending on their preferences.

In general, any of the constraints of optimization problem $(P_\theta[\mathcal{J}])$ might be superfluous and thus unnecessary in a certain MILP instance. Nevertheless, it is often the case that, because of the nature and structure of the MILP under consideration, there is some group of constraints, say $\tilde{\mathcal{J}} \subset \mathcal{J}$, which are more prone to be redundant and/or whose elimination from the original MILP brings a substantial reduction in computational time. It may be, therefore, very useful to focus on that group of constraints only, and adding the set of constraints $\mathcal{J} \setminus \tilde{\mathcal{J}}$ into \mathcal{S}_i and \mathcal{B}_t , $\forall t$, by default. We notice that this is an issue analogous to that of deciding which group of binary variables is better to be learned in those strategies that help solve MILPs very fast by predicting the optimal value of some of these variables, see, for instance, [19]. In the numerical experiments we present in Section 4.2, we will specify which constraints are considered in $\tilde{\mathcal{J}}$.

The efficiency of our proposal is measured using different performance metrics depending on the size of the datasets. For the toy example of Section 4.2.1, we show the behavior of our proposal \mathcal{S} -learner + CG compared with the alternative machine-learning-aided approach \mathcal{B} -learner + CG. We compare both strategies in terms of: (i) the constraints used to warm-start the CG method, (ii) the number of runs of CG, and (iii) the final reduced set of constraints needed to get the optimal solution of the MILP. In contrast, for the large datasets of Sections 4.2.2 and 4.2.3, we measure the benefits of our approach over the T runs in terms of: (i) the minimum and the maximum number of constraints considered in the reduced MILPs. Such a number of constraints is denoted as C_{min} , and C_{max} , respectively; (ii) the minimum and maximum number of iterations executed in the constraint generation procedure of the three methodologies averaged over all test instances, denoted by I_{min} and I_{max} , respectively; (iii) the percentage, P_1 , of instances that require only one iteration of the constraint generation method; and (iv) the percentage of online computational burden in comparison with the original MILP formulation, defined as $\Delta = \frac{1}{T} \sum_{t=1}^T \delta_t$, where $\delta_t = 100 \frac{\tau_t^{pred} + \tau_t^{CG}}{\tau_t^{MILP}}$. For each instance t , τ_t^{MILP} denotes the time needed to solve the original MILP with the whole set of constraints \mathcal{J} . In addition, τ_t^{pred} is the time employed in predicting sets \mathcal{B}_i or \mathcal{S}_i by the knn strategy. Finally, the notation τ_t^{CG} indicates the computational time employed in all the iterations of the constraint generation method.

Apart from providing values that summarize the good performance of our approach, we have included three figures to illustrate the distribution of the performance measures. In particular, we provide the boxplots of the performance values obtained after running T iterations for the different methodologies in terms of: (i) the number of constraints considered in the reduced MILPs, (ii) the number of iterations executed by the constraint generation strategy, and (iii) the percentage of the computational time spent online compared to the original MILP formulation, δ_t .

All the experiments have been carried out on a cluster with 21 Tb RAM, running Suse Leap 42 Linux distribution. MILP problems

are coded in Python 3.8 and Pyomo 6.1.2 and solved using Cplex 20.1.0. All the solver tuning parameters have been set to their default value except the *mixed integer optimality gap tolerance* (*mipgap*) that has been fixed to $1e^{-10}$. Finally, in order to make our approach transparent, we have saved the data and the code of our proposal in [20].

4.2. Case studies

The proposed methodology has been tested on three case studies: A toy example (Section 4.2.1), a synthetic MILP (Section 4.2.2) and a real-world application, the Unit Commitment Problem (Section 4.2.3).

4.2.1. Toy example

This section presents a toy example to illustrate how our proposal works. To this aim, we formulate the optimization problem in (2) with two decision variables, namely, $x \in \mathbb{R}$ and $y \in \mathbb{Z}$, and the feasible region given by the six constraints defined in (2b)–(2g).

$$\begin{cases} \min_{\substack{x \in \mathbb{R}, \\ y \in \mathbb{Z}}} x - y & (2a) \\ \text{s.t. } x \leq 1.5 & (2b) \\ y \leq 1.75 & (2c) \\ x \geq 0.5 & (2d) \\ x + y \geq b & (2e) \\ y \geq 0 & (2f) \\ y \leq 2.25 & (2g) \end{cases}$$

Note that problem (2) depends on the parameter $b \in \mathbb{R}$ that appears in the right-hand side of constraint (2e). Hence, in this toy example, we have $\theta = b$. We assume given a database containing the results of three optimization problems solved for the values $b \in \{1, 1.25, 1.5\}$. Fig. 2 shows the feasible region of the problem for the three different values of b . An arrow at the bottom right corner indicating the direction of improvement of the objective function is also depicted. It is easy to see that the optimal solution for these three problems is the point $A = (0.5, 1)$.

Now, let us assume that we have a new unobserved problem instance given by the value $b^{test} = 1.3$. The objective of this section is to learn a reduced set of constraints that will be used to initialize the constraint generation method and solve the problem for b^{test} . This way, the number of iterations needed to find the optimal solution of the MILP is reduced, and so does the associated computational burden. In this toy example, we illustrate that it is not enough to learn such a reduced set of constraints with the information given by the binding constraints set, \mathcal{B} . We have to take into account also some crucial non-binding constraints included in the invariant constraint set \mathcal{S} obtained after running Algorithm 1. Table 1 shows which are the constraints that belong to both sets \mathcal{B} and \mathcal{S} for the three training instances given by values $b \in \{1, 1.25, 1.5\}$. Note that, as stated in Section 4.1, the constraints belonging to the set \mathcal{B} (resp. \mathcal{S}) are labeled with 1. On the other hand, the constraints that do not belong to \mathcal{B} (resp. \mathcal{S}) are denoted with -1 .

The main difference between sets \mathcal{B} and \mathcal{S} is the constraint (2c), which is included in the invariant constraint set \mathcal{S} for the three values of b , but not in \mathcal{B} . Importantly, constraint (2c) is the unique non-binding constraint needed to recover the optimal solution given by the point A . In effect, if this constraint is




Fig. 2. Feasible region, direction of improvement of the objective function and optimal solution (point A) of the three instances obtained with the values $b \in \{1, 1.25, 1.5\}$ in Problem (2).

Table 1
Constraints included in the sets \mathcal{B} and \mathcal{S} for the values $b \in \{1, 1.25, 1.5\}$ of the three training instances.

	b	(2b)	(2c)	(2d)	(2e)	(2f)	(2g)
\mathcal{B}	1	-1	-1	1	-1	-1	-1
	1.25	-1	-1	1	-1	-1	-1
	1.5	-1	-1	1	1	-1	-1
\mathcal{S}	1	-1	1	1	-1	-1	-1
	1.25	-1	1	1	-1	-1	-1
	1.5	-1	1	1	1	-1	-1

Table 2
Performance results: Toy example.

	k	Warm-start constraint set	Iterations CG	Final set of constraints
\mathcal{B} -learner+CG	1	(2d)	2	(2c), (2d)
	2	(2d), (2e)	2	(2c), (2d), (2e)
	3	(2d), (2e)	2	(2c), (2d), (2e)
\mathcal{S} -learner+CG	1	(2c), (2d)	1	(2c), (2d)
	2	(2c), (2d), (2e)	1	(2c), (2d), (2e)
	3	(2c), (2d), (2e)	1	(2c), (2d), (2e)

removed, then the optimal solution moves from point A to point $B = (0.5, 2)$.

To find the initial set of constraints warm-starting the constraint generation method for the test problem instance given by $b^{test} = 1.3$, we train knn for $k \in \{1, 2, 3\}$ using the information provided by sets \mathcal{B} and \mathcal{S} in Table 1. Note that for $k = 1$, the closest problem instance to b^{test} is the one associated to $b = 1.25$. In addition, for $k = 2$, the closest neighbors are the MILPs given by $b = 1.25$ and $b = 1.5$. Finally, when $k = 3$, the three values of $b \in \{1, 1.25, 1.5\}$ are considered as the nearest neighbors of b^{test} . We collect in Table 2 the results obtained for both approaches in terms of: (i) the constraints initially selected, (ii) the number of iterations executed by the constraint generation strategy, and (iii) the constraints employed to solve the reduced test MILP instance. Due to the small size of the problem, the computational burden of both approaches is negligible. Consequently, this information has been omitted.

Several conclusions can be derived from the results shown in Table 2. Since we are executing a modified version of the

constraint generation method, it is guaranteed that the optimal solution of the problem instance given by b^{test} is reached using both data-driven approaches. Indeed, such an optimal solution is also attained at point A using a reduced set of two or three constraints (see the last column of Table 2). This means that we have managed to decrease between 50%–66% the cardinality of the original set of constraints, depending on the value of k we use. However, the number of iterations that the CG method needs to perform to get this optimum varies depending on which of the two data-driven approaches is employed. Indeed, the number of CG iterations executed by the approach \mathcal{S} -learner + CG is smaller than those employed in \mathcal{B} -learner + CG for all values of k (1 versus 2 iterations). This is due to the fact that the latter is trained just using the information taken from the binding constraints, which is not enough when there exist integer decision variables in the optimization problem. It is important to highlight that the approach \mathcal{B} -learner + CG is not able to find the optimal solution executing just one iteration of the CG method *even if all the available training instances are used*, that is, even running the knn method with $k = 3$. In contrast, our proposal \mathcal{S} -learner + CG always finds the optimal solution of the problem by only running the CG strategy once. This occurs because the non-binding constraint (2c) is included in the initial set of constraints used to warm-start the constraint generation method. Finally, as it was explained in Section 4.1, we want to be on the conservative side when choosing the knn voting choice. This is the reason why constraint (2e) is included in the initial set of constraints of both approaches for $k = 2$ and $k = 3$, even if such a constraint is not necessary to reach the optimal solution of the test problem instance. However, including such a small number of non-critical constraints have minor consequences in terms of computational times, as will be observed in the larger datasets of Sections 4.2.2 and 4.2.3.

4.2.2. Synthetic setup

In this section, we restrict ourselves to the MILP problem of the form (3):

$$\left\{ \begin{array}{l} \min_{x \in \mathbb{R}^n, y \in \{0, 1\}^n} \sum_{i=1}^n c_i x_i \quad (3a) \\ \text{s.t.} \quad \sum_{i=1}^n a_{ij} x_i \leq b_j, \quad j = 1, \dots, m \quad (3b) \\ l_i y_i \leq x_i \leq u_i y_i, \quad i = 1, \dots, n \quad (3c) \end{array} \right.$$

Table 3
Performance results: Synthetic MILP.

	k	$[C_{min}, C_{max}]$	$[I_{min}, I_{max}]$	$P_1(\%)$	$\Delta(\%)$
CG	–	[119, 132]	[120, 133]	0.0	1956.20
\mathcal{B} -learner+CG	1	[118, 129]	[12, 26]	0.0	1050.89
	5	[120, 130]	[1, 14]	2.8	366.71
	10	[122, 131]	[1, 9]	20.2	204.87
	50	[125, 132]	[1, 5]	58.4	119.76
	100	[127, 134]	[1, 4]	66.4	112.33
\mathcal{S} -learner+CG	500	[130, 135]	[1, 3]	83.8	87.90
	999	[134, 136]	[1, 3]	85.6	87.55
	1	[120, 130]	[1, 8]	25.0	195.67
	5	[124, 134]	[1, 4]	77.7	95.25
	10	[126, 134]	[1, 3]	90.6	81.06
\mathcal{S} -learner+CG	50	[131, 138]	[1, 3]	98.3	73.25
	100	[132, 139]	[1, 2]	99.3	74.47
	500	[137, 139]	[1, 1]	100.0	72.22
	999	[138, 139]	[1, 1]	100.0	74.38

where $\mathbf{a}_i = (a_{i1}, \dots, a_{im})^T, \forall i \leq n$, and $\mathbf{b} = (b_1, \dots, b_m)^T$ are column vectors in \mathbb{R}^m , and $\mathbf{c} = (c_1, \dots, c_n)^T, \mathbf{l} = (l_1, \dots, l_n)^T$ and $\mathbf{u} = (u_1, \dots, u_n)^T$ are column vectors in \mathbb{R}^n .

MILPs like (3) can be interpreted as linear programs where some of the continuous variables x_i have a forbidden zone within the range $(0, l_i)$. Consequently, problems like (3) contain the so-called logical constraints, where a continuous variable vanishes if the associated binary variable is zero. This type of problems, with a logical relationship between continuous and binary variables, has a wide variety of applications, as the authors of [21] explain. For a real-life example, one can think in a nuclear energy context. For instance, a nuclear unit whose maximum power is 1000 MW cannot generate energy within the range $[0, 500]$ due to the nuclear reactor stability.


As mentioned in Section 4.1, it may be computationally productive to screen out only a subset of constraints $\tilde{\mathcal{J}} \subset \mathcal{J}$. In the case of Problem (3), for example, most of the constraints (3c) are expected to be binding at the optimal solution. Therefore, we consider in this example that $\tilde{\mathcal{J}}$ is solely formed by the m constraints in (3b). Consequently, the n constraints in (3c) are included by default into \mathcal{S}_t and $\mathcal{B}_t, \forall t$.

We assume given a database with $T = 1000$ optimization problems of the type of (3). Each optimization problem t comprises $m = 250$ constraints, $n = 500$ continuous variables, and $n = 500$ binaries. We assume that the problems just depend on the parameter $\theta = \mathbf{b}$, i.e., $\mathbf{a}_i, \forall i, \mathbf{c}, \mathbf{l}$, and \mathbf{u} remain fixed for the 1000 optimization problems.


To synthetically generate the database, the values for the parameters $\mathbf{a}_i, \forall i, \mathbf{c}, \mathbf{l}$, and \mathbf{u} (which the 1000 MILPs share) have been randomly selected according to a normal distribution with mean 0 and standard deviation 10, i.e., $\mathcal{N}(0, 10)$. Note that we assure that lower bounds l_i take on smaller values than the upper bounds u_i , i.e., $l_i < u_i, i = 1, \dots, n$. Then, 1000 parameter vectors \mathbf{b} have been generated again according to the same distribution $\mathcal{N}(0, 10)$. The entire database can be downloaded from [20].

Fig. 3 and Table 3 shows the performance metrics for the three methodologies. The data-driven strategies \mathcal{B} -learner + CG and \mathcal{S} -learner + CG include output results after running the knn algorithm for different values of $k \in \{1, 5, 10, 50, 100, 500, 999\}$. Training the knn algorithm with $k = 999$ is equivalent to running a naive method that includes a constraint j in the set \mathcal{S}_t (resp. \mathcal{B}_t) if that constraint is contained in at least one of the training sets $\mathcal{S}_t, \forall t$ (resp. $\mathcal{B}_t, \forall t$). Note also that due to their nature, it is obvious that the three algorithms recover the optimal solution of the original MILP instances.

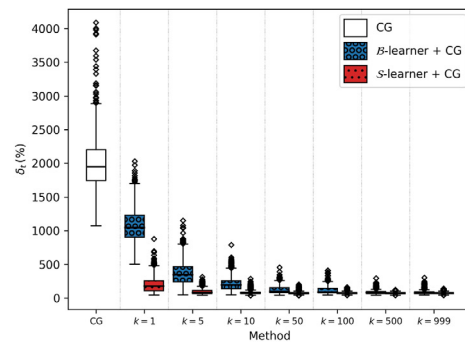
It can be observed that the number of iterations that need to be executed by the optimization-based method CG is significantly



(a) Number of constraints considered in the reduced MILPs.



(b) Number of CG iterations.



(c) Percentage of online computational burden in comparison with the original MILP formulation.

Fig. 3. Performance results: Synthetic MILP.

larger than those required by the data-driven methods \mathcal{B} -learner + CG and \mathcal{S} -learner + CG. Indeed, the y-axis of Fig. 3(b) has been divided into two different parts. This way, the large number of iterations needed in the pure optimization-based method CG does not affect the visualization of the number of CG iterations in the data-aided approaches. In addition, such a difference in the number of iterations clearly shows the benefits of using machine learning tools to alleviate the online computational burden in contrast with the use of pure optimization-assisted strategies.

Moreover, if we compare the number of CG iterations in both data-driven methods, \mathcal{B} -learner + CG and \mathcal{S} -learner + CG, we can see that the number of iterations to be run in the former approach is larger than in the latter, for all the values of k . Particularly, if we focus on the value $k = 1$ of Fig. 3(b), we observe that the minimum number of iterations that the method \mathcal{B} -learner + CG executed is larger than the maximum number of iterations of the strategy \mathcal{S} -learner + CG. Indeed, we can affirm that in order to reduce the number of iterations of the constraint generation

method, it is important to find a good initial set of constraints for the warm-start. Actually, if such an initial set is not good enough, solving the original MILP to optimality with the whole set of constraints, \mathcal{J} , could be better in terms of the computational burden than running a few iterations of the reduced MILP instances. For example, observe the output results of Table 3 of $k = 100$ for both approaches \mathcal{B} -learner + CG and \mathcal{S} -learner + CG. It can be seen that, in the first case, the computational load is increased around 12% on average, whereas in the second case, a reduction of approximately 25 percentage points is attained. In this regard, we should also highlight from Table 3 that running \mathcal{B} -learner + CG with values $k \in \{1, 5, 10, 50, 100\}$ provides no online running time benefits. In contrast, online computational savings can be observed from $k = 5$ if \mathcal{S} -learner + CG is performed. That is, the use of the sets \mathcal{S}_t , $\forall t$ which are built offline, is advantageous from a computational point of view.

In addition, we emphasize the results from Table 3 of $k = 500$ and $k = 999$ for the \mathcal{S} -learner + CG. Note that the number of CG iterations in both cases is exactly one. That means that the prediction of the invariant constraints sets includes all the binding and non-binding constraints necessary to recover the optimal solution of the original MILP formulation. It is important to remark that the \mathcal{B} -learner + CG is unable to reproduce these results even for $k = 999$, i.e., even using a naive method with all the available data. The reason for this issue is that the strategy \mathcal{B} -learner + CG is only based on the binding constraints sets, \mathcal{B}_t , $\forall t$, which are not sufficient to recover the optimal solution. Therefore, the risk of generating reduced MILPs, which are not equivalent to the original ones, increases.

Regarding the number of constraints of the reduced MILP formulations, it can be observed in Table 3 and Fig. 3(a) that the number of constraints necessary for solving the unseen instances includes around 50% – 55% of the total number of constraints, independently of which of the three algorithms is run. This means that the size of the original optimization problem is considerably reduced with any of the methodologies. Regarding the data-driven methods, the number of retained constraints in the \mathcal{S} -based methods is, as expected, slightly greater than in the \mathcal{B} -based ones. The difference boils down to a few extra constraints, which, thus, barely increases the size of the reduced MILPs. Very importantly, however, leaving these few constraints in the reduced MILPs has a major impact on the number of CG iterations, and thus in the online computational load, as Table 3 and Fig. 3(c) show.

In this vein, it is essential to remark that the simple fact of including any type of constraint in the set \mathcal{S}_t is not enough to recover the optimal solution in a shorter computational time. For instance, notice in Fig. 3(a) that the distribution of retained constraints after training the \mathcal{B} -learner + CG with $k = 100$, on the one hand, and the \mathcal{S} -learner + CG with $k = 10$, on the other, is very similar. However, looking at Table 3 we can see that there is no online time reduction in the first case, whereas around 80% of the computational load is employed in the second case. This is because more critical non-binding constraints are retained in the reduced MILPs in the latter case.

To sum up, our strategy is able to attain the optimal solution of an MILP, thanks to the efficient warm-start of the constraint generation method. Actually, we are able to correctly identify offline an invariant constraint set of an MILP, \mathcal{S}_t , $\forall t$. In doing so, the performance of the machine learning tool (knn in our case) that is used to initialize the constraint generation method is improved. This improvement substantially decreases the number of iterations performed at the expense of a slight increment in the cardinality of the set of retained constraints. This increment, however, does not involve an increase in the online solution time. In addition, the resulting reduced problems are easier to solve than the original MILPs.

4.2.3. Real-world application: Unit Commitment problem

The Unit Commitment problem (UC) is one of the most important problems in power systems, as the authors in [12] affirm. The goal of UC is to determine, at minimum cost, the on/off status and the power to be dispatched by each generation unit in order to satisfy the electric demand. Mathematically, the (DC version of the) UC problem can be formulated as the following MILP:

$$\left\{ \begin{array}{l} \min_{x \in \mathbb{R}^n, y \in \{0, 1\}^n} \sum_{i=1}^n c_i x_i \quad (4a) \\ \text{s.t.} \quad \sum_{i=1}^n x_i = \sum_{i=1}^n d_i, \quad (4b) \\ -f_j \leq \sum_{i=1}^n a_{ij}(x_i - d_i) \leq f_j, \quad j = 1, \dots, m \quad (4c) \\ l_i y_i \leq x_i \leq u_i y_i, \quad i = 1, \dots, n \quad (4d) \end{array} \right.$$

where x_i is the power dispatched of generator i and y_i is a binary variable indicating whether the generator is turned on or turned off. In addition, c_i is the marginal cost of generator i , d_i is the electric demand at node i , a_{ij} are the so-called Power Transfer Distribution Factors (PTDF) in [12], and f_j , l_i and u_i are, respectively, flow and power generation limits. The objective function (4a) aims to minimize the total cost. Constraint (4b) is the power balance equation, enforcing the supply of the total demand in the power network. Constraints (4c) limit the flow of line j , given by $\sum_{i=1}^n a_{ij}(x_i - d_i)$, within the range $[-f_j, f_j]$. Finally, constraints (4d) ensures that the power dispatched x_i be within l_i and u_i if and only if generator i is turned on, i.e. if and only if $y_i = 1$. We remark that, for simplicity, formulation (4) considers that there is at most one generator connected to each network node.

The Unit Commitment problem is a suitable application to test the performance of our method for three reasons. First, the increasing integration of renewable sources in current power systems requires that the unit commitment problem be solved multiple times within short-time windows so that commitment decisions can be adapted to rapid changes in operating conditions. In practice, this means that this problem must be solved as fast as possible. Second, the unit commitment problem is solved several times per day with only minor changes in the input data. Therefore, historical data of previous instances are usually available to be used in learning tasks. Third, implementing commitment decisions that violate some of the security constraints of the UC may lead to catastrophic events such as power blackouts. Therefore, attaining optimal (and feasible) solutions is also a requirement for this practical application.


While the marginal cost and production limits of power plants and network topology do not typically change over a year, the electric demand suffers from daily and weekly fluctuations. Hence, we decide to fix parameters a_{ij} , l_i , u_i , f_j and c_i , $\forall i, j$, and just vary the input parameter $\theta = \mathbf{d}$. Moreover, it is known that, in practice, just a small percentage of the power flow constraints (4c) are binding at the optimum for typical power systems. See Ref. [22] for further details. Hence, we consider in this example that $\bar{\mathcal{J}}$ is made up of the $2m$ constraints in (4c). The entire dataset consists of $n = 96$ continuous and $n = 96$ binary variables and $m = 120$, leading to a total of 240 constraints (4c). The number of problem instances is $T = 8640$, corresponding to 360 days of data measured every hour. More details about this data can be found in [20].

The performance metrics of our methodology are collated in Table 4 and Fig. 4 for $k \in \{5, 10, 20, 50, 100\}$.


We can see from Fig. 4(a) that the reduced problems generated from the three algorithms have up to 85% fewer constraints than

Table 4
Performance results: Unit commitment.


	k	$[C_{min}, C_{max}]$	$[I_{min}, I_{max}]$	$P_1(\%)$	$\Delta(\%)$
CG	–	[0, 22]	[1, 23]	9.16	188.38
β -learner+CG	5	[0, 23]	[1, 8]	54.40	74.09
	10	[0, 25]	[1, 6]	62.01	63.19
	20	[0, 26]	[1, 5]	68.28	62.35
	50	[0, 27]	[1, 5]	76.90	54.56
	100	[0, 29]	[1, 5]	83.70	54.62
S -learner+CG	5	[0, 26]	[1, 5]	92.66	44.84
	10	[0, 28]	[1, 5]	97.21	40.57
	20	[0, 29]	[1, 4]	98.81	42.83
	50	[0, 30]	[1, 3]	99.45	40.57
	100	[0, 32]	[1, 3]	99.71	44.41



(a) Number of constraints considered in the reduced MILPs.



(b) Number of CG iterations.



(c) Percentage of online computational burden in comparison with the original MILP formulation.

Fig. 4. Performance results: Unit Commitment.

the original MILPs. However, such a reduction results in modest online computational savings when the β -learner + CG is used to train the knn classification model. In contrast, using the S -learner + CG manages to substantially improve the performance of the knn in terms of the online solution time. For instance, we can observe in Table 4 that the β -learner + CG with $k = 5$ employs around 75% of the time required to solve the original MILP formulation, whereas just 45% of the time is used if the knn is run on the S -learner + CG. This is of particular relevance in real-life applications, such as the UC problem, where the optimality guarantee in a short computational time is a priority.

Indeed, due to the good initialization of the constraint generation method, based on the invariant constraint set, our approach is able to get (almost always) a perfect prediction of the constraints needed to recover the optimal solution of the original MILP formulation, as the third column of Table 4 shows. Consequently, the number of iterations of the constraint generation method is 1 in most of the test instances, and the solution times are reduced, as Figs. 4(b) and 4(c) present. For instance, Table 4 states that our method found the optimal solution in approximately 93% of the instances if $k = 5$ is chosen. Moreover, if we observe the results for $k = 100$, nearly 100% of the problems attain the optimal MILP solution in one iteration of the constraint generation method.

To summarize, solving to optimality real-world optimization problems, such as the UC problem, may become a challenge. The presented results have shown that the warm-started constraint generation procedure proposed in this paper, which is based on the invariant constraint set, significantly reduces the computational burden and gets reduced MILPs that are equivalent to the original ones in terms of the optimal solution. Consequently, our approach strengthens and supports the use of machine-learning-aided optimization tools for online applications.

5. Conclusions and future work

Solving MILPs to optimality for online applications using traditional algorithms is not always possible due to their high computational burden. While the literature includes speed-up methods to solve MILP based on machine learning techniques, the obtained solutions may be suboptimal (or even infeasible in certain cases) for the original formulation. In this paper, we propose a machine-learning-aided warm-start constraint generation algorithm that attains the optimal solution of an MILP in a shorter computational time. The proposed approach is based on the offline identification of the invariant constraint sets of previous instances of the target MILP. In doing so, we significantly improve the prediction of invariant constraint sets for unseen instances. Thus, a much smaller number of iterations are needed to run the constraint generation algorithm, and the online computational burden is significantly diminished.

We compare our approach in a synthetic MILP and the unit commitment problem with the traditional constraint generation method and with a warm-started methodology that ignores the information given by the critical non-binding constraints. In both examples, the online computational time is significantly reduced with respect to the comparative strategies. For instance, in our experiments with the unit commitment application, the optimal solution is attained using around 40% of the time needed in the original MILP. This shows the advantage of using our methodology for solving MILP in online applications.

In our study, the MILP instances in the training set only differ on the right-hand side of their constraints. While our methodology can also be used with MILPs that are parameterized by the coefficient matrix, further investigation is required to evaluate its performance for this case. Another promising research line is the

design of methods for integrating information or expert knowledge on the MILP to be solved into the learning process in order to increase the predictive power of the machine learning engine on the MILP solution. Finally, our approach implies the training of an independent machine learning model for each constraint. In order to improve the performance metrics, a future research line would be to take advantage of the possible relationships among the constraints to train a unique machine learning algorithm.

CRedit authorship contribution statement

Asunción Jiménez-Cordero: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization, Funding acquisition. **Juan Miguel Morales:** Conceptualization, Methodology, Formal analysis, Investigation, Writing – original draft, Writing – review & editing, Visualization, Funding acquisition. **Salvador Pineda:** Conceptualization, Methodology, Formal analysis, Investigation, Writing – original draft, Writing – review & editing, Visualization, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported in part by the Spanish Ministry of Science and Innovation through project PID2020-115460GB-I00, in part by the European Research Council (ERC) under the EU Horizon 2020 research and innovation program (grant agreement No. 755705) in part, by the Junta de Andalucía (JA), the Universidad de Málaga (UMA), and the European Regional Development Fund (FEDER) through the research projects P20_00153 and UMA2018-FEDERJA-001, and in part by the Research Program for Young Talented Researchers of the University of Málaga under Project B1-2020-15. The authors thankfully acknowledge the computer resources, technical expertise, and assistance provided by the SCBI (Supercomputing and Bioinformatics) center of the University of Málaga. Funding for open access charge: Universidad de Málaga / CBUA.

References

- [1] C. Gambella, B. Ghaddar, J. Naoum-Sawaya, Optimization problems for machine learning: A survey, *European J. Oper. Res.* 290 (3) (2021) 807–828, <http://dx.doi.org/10.1016/j.ejor.2020.08.045>.

- [2] F. Lin, J. Wang, N. Zhang, J. Xiahou, N. McDonald, Multi-kernel learning for multivariate performance measures optimization, *Neural Comput. Appl.* 28 (8) (2017) 2075–2087, <http://dx.doi.org/10.1007/s00521-015-2164-9>.
- [3] L. Yang, A. Shami, On hyperparameter optimization of machine learning algorithms: Theory and practice, *Neurocomputing* 415 (2020) 295–316, <http://dx.doi.org/10.1016/j.neucom.2020.07.061>.
- [4] I.A. Zamfirache, R.-E. Precup, R.-C. Roman, E.M. Petriu, Policy iteration reinforcement learning-based control using a grey wolf optimizer algorithm, *Inform. Sci.* 585 (2022) 162–175, <http://dx.doi.org/10.1016/j.ins.2021.11.051>.
- [5] M. Conforti, G. Cornuejols, G. Zambelli, *Integer Programming*, Springer Publishing Company, Incorporated, 2014.
- [6] L.A. Wolsey, Mixed integer programming, in: *Wiley Encyclopedia of Computer Science and Engineering*, American Cancer Society, 2008, pp. 1–10, <http://dx.doi.org/10.1002/9780470050118.ecse244>.
- [7] Y. Bengio, A. Lodi, A. Prouvost, Machine learning for combinatorial optimization: A methodological tour d'horizon, *European J. Oper. Res.* 290 (2) (2021) 405–421, <http://dx.doi.org/10.1016/j.ejor.2020.07.063>.
- [8] D. Bertsimas, B. Stellato, The voice of optimization, *Mach. Learn.* 110 (2) (2021) 249–277, <http://dx.doi.org/10.1007/s10994-020-05893-5>.
- [9] A. Chmiela, E.B. Khalil, A. Gleixner, A. Lodi, S. Pokutta, Learning to schedule heuristics in branch and bound, in: A. Beygelzimer, Y. Dauphin, P. Liang, J.W. Vaughan (Eds.), *Advances in Neural Information Processing Systems*, 2021, URL: <https://openreview.net/forum?id=mvEhklqn45>.
- [10] M. Minoux, Networks synthesis and optimum network design problems: Models, solution methods and applications, *Networks* 19 (3) (1989) 313–360, <http://dx.doi.org/10.1002/net.3230190305>.
- [11] S. Pineda, J.M. Morales, A. Jiménez-Cordero, Data-driven screening of network constraints for unit commitment, *IEEE Trans. Power Syst.* 35 (5) (2020) 3695–3705, <http://dx.doi.org/10.1109/TPWRS.2020.2980212>.
- [12] A. Gómez-Exposito, A.J. Conejo, C. Cañizares, *Electric Energy Systems: Analysis and Operation*, CRC Press, 2018.
- [13] A.S. Xavier, F. Qiu, S. Ahmed, Learning to solve large-scale security-constrained unit commitment problems, *INFORMS J. Comput.* 33 (2) (2021) 739–756, <http://dx.doi.org/10.1287/ijoc.2020.0976>.
- [14] G.C. Calafiore, Random convex programs, *SIAM J. Optim.* 20 (6) (2010) 3427–3464, <http://dx.doi.org/10.1137/090773490>.
- [15] D. Bertsimas, J. Tsitsiklis, *Introduction to Linear Optimization*, first ed., Athena Scientific, 1997.
- [16] J. Friedman, T. Hastie, R. Tibshirani, *The Elements of Statistical Learning*, in: *Springer Series in Statistics*, vol. 1, Springer, Berlin, 2001.
- [17] T. Hastie, R. Tibshirani, J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer Science & Business Media, 2009.
- [18] K. Taunk, S. De, S. Verma, A. Swetapadma, A brief review of nearest neighbor algorithm for learning and classification, in: *2019 International Conference on Intelligent Computing and Control Systems*, ICCS, 2019, pp. 1255–1260, <http://dx.doi.org/10.1109/ICCS45141.2019.9065747>.
- [19] A. Lodi, L. Mossina, E. Rachelson, Learning to handle parameter perturbations in combinatorial optimization: An application to facility location, *EURO J. Transp. Logist.* 9 (4) (2020) 100023, <http://dx.doi.org/10.1016/j.ejtl.2020.100023>.
- [20] OASYS, *Warm_starting_CG_for_MIO_ML*, 2022, https://github.com/groupoasys/Warm_starting_CG_for_MIO_ML.
- [21] D. Bertsimas, R. Cory-Wright, J. Pauphilet, A unified approach to mixed-integer optimization problems with logical constraints, *SIAM J. Optim.* 31 (3) (2021) 2340–2367, <http://dx.doi.org/10.1137/20M1346778>.
- [22] F. Bouffard, F.D. Galiana, J.M. Arroyo, Umbrella contingencies in security-constrained optimal power flow, in: *15th Power Systems Computation Conference*, Vol. 5, PSCC, 2005.