

Learning to Project in Multi-Objective Binary Linear Programming

Alvaro Sierra-Altamiranda

Department of Industrial and Management System Engineering, University of South Florida, Tampa, FL, 33620 USA,
amsierra@mail.usf.edu, <http://www.eng.usf.edu/~amsierra/>

Hadi Charkhgard

Department of Industrial and Management System Engineering, University of South Florida, Tampa, FL, 33620 USA,
hcharkhgard@usf.edu, <http://www.eng.usf.edu/~hcharkhgard/>

Iman Dayarian

Culverhouse College of Business, The University of Alabama, Tuscaloosa, AL 35487 USA, idayarian@cba.ua.edu,
<https://culverhouse.ua.edu/news/directory/iman-dayarian/>

Ali Eshragh

School of Mathematical and Physical Sciences, The University of Newcastle, Callaghan, NSW 2308 Australia,
ali.eshragh@newcastle.edu.au, <https://www.newcastle.edu.au/profile/ali-eshragh>

Sorna Javadi

Department of Industrial and Management Systems Engineering, University of South Florida, Tampa, FL, 33620 USA,
javadis@mail.usf.edu

In this paper, we investigate the possibility of improving the performance of multi-objective optimization solution approaches using machine learning techniques. Specifically, we focus on multi-objective binary linear programs and employ one of the most effective and recently developed criterion space search algorithms, the so-called KSA, during our study. This algorithm computes all nondominated points of a problem with p objectives by searching on a projected criterion space, i.e., a $(p - 1)$ -dimensional criterion space. We present an effective and fast learning approach to identify on which projected space the KSA should work. We also present several generic features/variables that can be used in machine learning techniques for identifying the best projected space. Finally, we present an effective bi-objective optimization based heuristic for selecting the best subset of the features to overcome the issue of overfitting in learning. Through an extensive computational study over 2000 instances of tri-objective Knapsack and Assignment problems, we demonstrate that an improvement of up to 12% in time can be achieved by the proposed learning method compared to a random selection of the projected space.

Key words: Multi-objective optimization, machine learning, binary linear program, criterion space search algorithm, learning to project

History:

1. Introduction

Many real-life optimization problems involve multiple objective functions and they can be stated as follows:

$$\min_{\mathbf{x} \in \mathcal{X}} \{z_1(\mathbf{x}), \dots, z_p(\mathbf{x})\}, \quad (1)$$

where $\mathcal{X} \subseteq \mathbb{R}^n$ represents the set of feasible solutions of the problem, and $z_1(\mathbf{x}), \dots, z_p(\mathbf{x})$ are p objective functions. Because objectives are often competing in a multi-objective optimization problem, an ideal feasible solution that can optimize all objectives at the same time does not often exist in practice. Hence, when solving such a problem, the goal is often generating some (if not all) *efficient* solutions, i.e., a feasible solution in which it is impossible to improve the value of one objective without making the value of any other objective worse.

The focus of this study is on Multi-Objective Binary Linear Programs (MOBLPs), i.e., multi-objective optimization problems in which all decision variables are binary and all objective functions and constraints are linear. In the last few years, significant advances have been made in the development of effective algorithms for solving MOBLPs, see for instance Boland et al. (2015a,b, 2016, 2017b), Dächert et al. (2012), Dächert and Klamroth (2015), Fattahi and Turkay (2017), Kirlik and Sayın (2014), Özpeynirci and Köksalan (2010), Lokman and Köksalan (2013), Özlen et al. (2013), Przybylski and Gandibleux (2017), Przybylski et al. (2010), Soylu and Yıldız (2016), and Vincent et al. (2013). Many of the recently developed algorithms fall into the category of criterion space search algorithms, i.e., those that work in the space of objective functions' values. Hence, such algorithms are specifically designed to find all *nondominated points* of a multi-objective optimization problem, i.e., the image of an efficient solution in the criterion space is being referred to as a nondominated point. After computing each nondominated point, criterion space search algorithms remove the proportion of the criterion space dominated by the obtained nondominated point and search for not-yet-found nondominated points in the remaining space.

In general, to solve a multi-objective optimization problem, criterion space search algorithms solve a sequence of single-objective optimization problems. Specifically, when solving a problem with p objective functions, many criterion space search algorithms first attempt to transform the problem into a sequence of problems with $(p-1)$ objectives (Boland et al. 2017b). In other words, they attempt to compute all nondominated points by discovering

their projections in a $(p - 1)$ -dimensional criterion space. Evidently, the same process can be applied recursively until a sequence of single-objective optimization problems are generated. For example, to solve each problem with $(p - 1)$ objectives, a sequence of problems with $(p - 2)$ objectives can be solved.

Overall, there are at least two possible ways to apply the projection from a higher dimensional criterion space (for example p) to a criterion space with one less dimension (for example $p - 1$):

- *Weighted Sum Projection*: A typical approach used in the literature (see for instance Özlen and Azizoglu 2009) is to select one of the objective functions available in the higher dimension (for example $z_1(\mathbf{x})$) and remove it after adding it with some strictly positive weight to the other objective functions. In this case, by imposing different bounds for $z_1(\mathbf{x})$ and/or the value of the other objective functions, a sequence of optimization problems with $p - 1$ objectives will be generated.

- *Lexicographical Projection*: We first note that a *lexicographical optimization problem* is a two-stage optimization problem that attempts to optimize a set of objectives, the so-called *secondary* objectives, over the set of solutions that are optimal for another objective, the so-called *primary* objective. The first stage in the lexicographical optimization problem is a single-objective optimization problem as it optimizes the primary goal. The second stage, however, can be a multi-objective optimization problem as it optimizes the secondary objectives. Based on this definition, another typical approach (see for instance Özlen et al. (2013)) for projection is to select one of the objective functions available in the higher dimension (for example $z_1(\mathbf{x})$) and simply remove it. In this case, by imposing different bounds for $z_1(\mathbf{x})$ and/or the value of the other objective functions, a sequence of lexicographical optimization problems should be solved in which $z_1(x)$ is the primary objective and the remaining $p - 1$ objectives are secondary objectives.

In light of the above, which objective function should be selected for doing a projection and how to do a projection are two typical questions that can be asked when developing a criterion space search algorithm. So, by this observation, there are many possible ways to develop a criterion space search algorithm and some of which may perform better for some instances. So, the underlying research question of this study is that whether Machine Learning (ML) techniques can help us answer the above questions for a given class of instances of a multi-objective objective optimization problem?

It is worth mentioning that, in recent years, similar questions have been asked in the field of single-objective optimization. For example, ML techniques have been successfully implemented for the purpose of *variable selection* and *node selection* in branch-and-bound algorithms (see for instance Khalil et al. (2016), Alvarez et al. (2017), Sabharwal et al. (2012), He et al. (2014), Khalil et al. (2017)). However, still the majority of the algorithmic/theoretical studies in the field of ML have been focused on using optimization models and algorithms to enhance ML techniques and not the other way around (see for instance Roth and Yih (2005), Bottou (2010), Le et al. (2011), Sra et al. (2012), Snoek et al. (2012), Bertsimas et al. (2016)).

In general, to the best of our knowledge, there are no studies in the literature that address the problem of enhancing multi-objective optimization algorithms using ML. In this study, as the first attempt, we focus only on the simplest and most high-level question that can be asked, that is for a given instance of MOBLP with p objective functions which objective should be removed for reducing the dimension of the criterion space to $p - 1$ in order to minimize the solution time?

It is evident that if one can show that ML is even valuable for such a high-level question then deeper questions can be asked and explored that can possibly improve the solution time significantly. In order to answer the above question, we employ one of the effective state-of-the-art algorithms in the literature of multi-objective optimization, the so-called KSA which is developed by Kirlik and Sayın (2014). This algorithm uses the lexicographical projection for reducing the p -dimensional criterion space to $p - 1$, and then it recursively reduces the dimension from $p - 1$ to 1 by using a special case of the weighted sum projection in which all the weights are equal to one.

Currently, the default objective function for conducting the projection from p -dimensional criterion space to $p - 1$ is the first objective function (or better say random because one can change the order of the objective functions in an input file). So, a natural question is that does it really matter which objective function is selected for such a projection? To answer this question, we conducted a set of experiments by using a C++ implementation of the KSA which is publicly available in <http://home.ku.edu.tr/~moolibrary> and recorded the number of ILPs solved (#ILPs) and the computational time (in seconds). We generated 1000 instances (200 per class) of tri-objective Assignment Problem (AP) and 1000 instances (200 per class) of Knapsack Problem (KP) based on the

procedure described by Kirlik and Sayın (2014). Table 1 shows the impact of projecting based on the worst and best objective function using the KSA where #ILPs is the number of single-objective integer linear programs solved. Numbers reported in this table are averages over 200 instances.

Table 1 Projecting based on different objectives using the KSA.

Type	#Objectives	#Variables	Projecting worst objective		Projecting best objective		%Decrease	
			Run time (s.)	#ILPs	Run time (s.)	#ILPs	Run time (s.)	#ILPs
AP	3	20 × 20	351.56	5,122.67	337.03	5,015.39	4.31%	4.61%
		25 × 25	948.21	9,685.69	912.07	9,500.52	3.96%	4.40%
		30 × 30	2,064.34	16,294.37	1,988.64	16,019.64	3.81%	4.01%
		35 × 35	4,212.69	26,161.34	4,050.44	25,592.13	4.01%	4.27%
		40 × 40	6,888.45	35,737.21	6,636.79	35,061.41	3.79%	3.97%
KP	3	60	270.34	3,883.45	212.41	3,861.68	27.27%	1.05%
		70	813.31	6,182.04	638.87	6,158.20	27.31%	0.82%
		80	1,740.80	9,297.96	1,375.58	9,265.09	26.55%	0.74%
		90	5,109.56	14,257.32	3,917.32	14,212.35	30.44%	0.63%
		100	10,451.97	19,420.06	7,780.96	19,366.88	34.33%	0.57%

We observe that, on average, the running time can be reduced up to 34% while the #ILPs can be improved up to 4%. This numerical study clearly shows the importance of projection in the solution time. Hence, it is certainly worth studying ML techniques in predicting the best objective function for projecting, the so-called *learning to project*. So, our main contribution in this study is to introduce an ML framework to simulate the selection of the best objective function to project. We collect data from each objective function and their interactions with the decision space to create features. Based on the created features, an easy-to-evaluate function is learned to emulate the classification of the projections. Another contribution of this study is developing a simple but effective bi-objective optimization-based heuristic approach to select the best subset of features to overcome the issue of overfitting. We show that the accuracy of the proposed prediction model can reach up to around 72%, which represents up to 12% improvement in solution time.

The rest of this paper is organized as follows. In Section 2, some useful concepts and notations about multi-objective optimization are introduced and also a high-level description of the KSA is given. In Section 3, we provide a high-level description of our proposed

machine learning framework and its three main components. In Section 4, the first component of the framework, which is a pre-ordering approach for changing the order of the objective functions in an input file, is explained. In Section 5, the second component of the framework that includes features and labels are explained. In Section 6, the third/last component of the framework, which is a bi-objective optimization based heuristic for selecting the best subset of features, is introduced. In Section 7, we provide a comprehensive computational study. Finally, in Section 8, we provide some concluding remarks.

2. Preliminaries

A *Multi-Objective Binary Linear Program* (MOBLP) is a problem of the form (1) in which $\mathcal{X} := \{\mathbf{x} \in \{0, 1\}^n : A\mathbf{x} \leq \mathbf{b}\}$ represents the *feasible set in the decision space*, $A \in \mathbb{R}^{m \times n}$, and $\mathbf{b} \in \mathbb{R}^m$. It is assumed that \mathcal{X} is *bounded* and $z_i(\mathbf{x}) = \mathbf{c}_i^\top \mathbf{x}$ where $\mathbf{c}_i \in \mathbb{R}^n$ for $i = 1, 2, \dots, p$ represents a linear objective function. The image \mathcal{Y} of \mathcal{X} under vector-valued function $\mathbf{z} := (z_1, z_2, \dots, z_p)^\top$ represents the *feasible set in the objective/criterion space*, that is $\mathcal{Y} := \{\mathbf{o} \in \mathbb{R}^p : \mathbf{o} = \mathbf{z}(\mathbf{x}) \text{ for all } \mathbf{x} \in \mathcal{X}\}$. Throughout this article, vectors are always column-vectors and denoted in bold font.

DEFINITION 1. A feasible solution $\mathbf{x} \in \mathcal{X}$ is called *efficient* or *Pareto optimal*, if there is no other $\mathbf{x}' \in \mathcal{X}$ such that $z_i(\mathbf{x}') \leq z_i(\mathbf{x})$ for $i = 1, \dots, p$ and $\mathbf{z}(\mathbf{x}') \neq \mathbf{z}(\mathbf{x})$. If \mathbf{x} is efficient, then $\mathbf{z}(\mathbf{x})$ is called a *nondominated point*. The set of all efficient solutions $\mathbf{x} \in \mathcal{X}$ is denoted by \mathcal{X}_E . The set of all nondominated points $\mathbf{z}(\mathbf{x}) \in \mathcal{Y}$ for some $\mathbf{x} \in \mathcal{X}_E$ is denoted by \mathcal{Y}_N and referred to as the *nondominated frontier*.

Overall, multi-objective optimization is concerned with finding all nondominated points, i.e., an exact representation of the elements of \mathcal{Y}_N . The set of nondominated points of a MOBLPs is finite (since by assumption \mathcal{X} is bounded). However, due to the existence of *unsupported* nondominated points, i.e., those nondominated points that cannot be obtained by optimizing any positive weighted summation of the objective functions over the feasible set, computing all nondominated points is challenging. One of the effective criterion space search algorithms for MOBLPs is the KSA and its high-level description is provided next.

The KSA is basically a variation of the ε -constraint method for generating the entire nondominated frontier of multi-objective integer linear programs. In each iteration, this algorithm solves the following lexicographical optimization problem in which the first stage is:

$$\hat{\mathbf{x}} \in \arg \min \{z_1(\mathbf{x}) : \mathbf{x} \in \mathcal{X}, z_i(\mathbf{x}) \leq u_i \quad \forall i \in \{2, \dots, p\}\},$$

where u_2, \dots, u_p are user-defined upper bounds. If $\hat{\mathbf{x}}$ exists, i.e., the first stage is feasible, then the following second-stage problem will be solved:

$$\hat{\mathbf{x}}^* \in \arg \min \left\{ \sum_{i=2}^p z_i(\mathbf{x}) : \mathbf{x} \in \mathcal{X}, z_1(\mathbf{x}) \leq z_1(\hat{\mathbf{x}}), z_i(\mathbf{x}) \leq u_i \quad \forall i \in \{2, \dots, p\} \right\}.$$

The algorithm computes all nondominated points by imposing different values on u_2, \dots, u_p in each iteration. Interested readers may refer to Kirlik and Sayın (2014) for further details about how values of u_2, \dots, u_p will be updated in each iteration. It is important to note that in the first stage users can replace the objective function $z_1(\mathbf{x})$ with any other arbitrary objective function, i.e., $z_j(\mathbf{x})$ where $j \in \{1, \dots, p\}$, and change the objective function of the second stage accordingly, i.e., $\sum_{i=1: i \neq j}^p z_i(\mathbf{x})$. As shown in Introduction, on average, the running time can decrease up to 34% by choosing the right objective function for the first stage. So, the goal of the proposed machine learning technique in this study is to identify the best choice.

As an aside, we note that to be consistent with our explanation of the lexicographic and/or weighted sum projections in Introduction, the lexicographic optimization problem of the KSA is presented slightly differently in this section. Specifically, Kirlik and Sayın (2014) use the following optimization problem instead of the second-stage problem (mentioned above):

$$\hat{\mathbf{x}}^* \in \arg \min \left\{ \sum_{i=1}^p z_i(\mathbf{x}) : \mathbf{x} \in \mathcal{X}, z_1(\mathbf{x}) = z_1(\hat{\mathbf{x}}), z_i(\mathbf{x}) \leq u_i \quad \forall i \in \{2, \dots, p\} \right\}.$$

However, one can easily observe that these two formulations are equivalent. In other words, the lexicographic optimization problem introduced in this section is a just different representation of the one proposed by Kirlik and Sayın (2014).

3. Machine learning framework

We now introduce our ML framework for learning to project in MOBLPs. Our proposed framework is based on Multi-class Support Vector Machine (MSVM). In this application, MSVM learns a function $f : \Phi \rightarrow \Omega$ from a training set to predict which objective function will have the best performance in the first stage of the KSA (for a MOBLP instance),

where Φ is the feature map domain describing the MOBLP instance and $\Omega := \{1, 2, \dots, p\}$ is the domain of the labels. A label $y \in \Omega$ indicates the index of the objective function that should be selected. We do not explain MSVM in this study but interested readers may refer to Crammer and Singer (2001) and Tsochantaridis et al. (2004) for details. We used the publicly available implementation of MSVM in this study which can be found in <https://goo.gl/4hLjyq>. It is worth mentioning that we have used MSVM mainly because it was performing well during the course of this study. In Section 7.4, we also report results obtained by replacing MSVM with Random Forest (Breiman 2001, Prinzie and Van den Poel 2008) to show the performance of another learning technique in our proposed framework. Also, we provide more reasons in Section 7.4 about why MSVM is used in this study. Overall, the proposed ML framework contains three main components:

- *Component 1:* It is evident that by changing the order of the objective functions of a MOBLP instance in an input file, the instance remains the same. Therefore, in order to increase the stability of the prediction of MSVM, we propose an approach to pre-order the objective functions of each MOBLP instance in an input file before feeding it to MSVM (see Section 4).

- *Component 2:* We propose several generic features that can be used to describe each MOBLP instance. A high-level description of the features can be found in Section 5 and their detailed descriptions can be found in Appendix A.

- *Component 3:* We propose a bi-objective heuristic approach (see Section 6) for selecting the best subset of features for each class of MOBLP instances (which are AP and KP in this study). Our numerical results show that our approach selects around 15% of features based on the training set for each class of MOBLP instances in practice. Note that identifying the best subset of features is helpful for overcoming the issue of overfitting and improving the prediction accuracy (Charkhgard and Eshragh 2019, Tibshirani 1996).

The proposed ML framework uses the above components for training purposes. A detailed discussion on the accuracy of the proposed framework on a testing set (for each class of MOBLP instances) is given in Section 7.

4. A pre-ordering approach for objective functions

It is obvious that by changing the order of objective functions in an input file corresponding to an instance, a new instance will not be generated. In other words, only the instance is represented differently in that case and hence its nondominated frontier will remain the

same. This suggests that the vector of features that will be extracted for any instance should be independent of the order of the objective functions. To address this issue, we propose to perform a pre-ordering (heuristic) approach before giving an instance to MSVM for training or testing purposes. That is, when users provide an instance, we first change its input file by re-ordering the objective functions before feeding it to the MSVM. Obviously, this somehow stabilizes the prediction accuracy of the proposed ML framework.

In light of the above, let

$$\tilde{\mathbf{x}} := \left(\frac{1}{\sum_{i=1}^p |c_{i1}| + 1}, \dots, \frac{1}{\sum_{i=1}^p |c_{in}| + 1} \right).$$

In the proposed approach, we re-order the objective functions in an input file in a non-decreasing order of $\mathbf{c}_1^\top \tilde{\mathbf{x}}, \mathbf{c}_2^\top \tilde{\mathbf{x}}, \dots, \mathbf{c}_p^\top \tilde{\mathbf{x}}$. Intuitively, $\mathbf{c}_i^\top \tilde{\mathbf{x}}$ can be viewed as the normalization score for objective function $i \in \{1, \dots, p\}$. In the rest of this paper, the vector \mathbf{c}_i for $i = 1, 2, \dots, p$ is assumed to be ordered according to the proposed approach.

5. Features and label describing a MOBLP instance

This section provides a high-level explanation of the features that we create to describe a MOBLP instance and also how each instance is labeled. To the best of our knowledge, there are no studies that introduce features to describe multi-objective instances, and hence the proposed features are new.

5.1. Features

The efficiency of our proposed ML approach relies on the features describing a MOBLP instance. In other words, the features should be easy to compute and effective. Based on this observation, we create only *static* features, i.e., those that are computed once using just the information provided by the MOBLP instance. Note that we only consider static features because the learning process and the decision on which objective function to select for projection (in the KSA) have to take place before solving the MOBLP instance. Overall, due to the nature of our research, most of our features describe the objective functions of the instances. We understand that the objective functions by themselves are a limited source of information to describe an instance. Therefore, we also consider establishing some relationships between the objective functions and the other characteristics of the instance in order to improve the reliability of our features.

In light of the above, a total of $5p^2 + 106p - 50$ features are introduced for describing each instance of MOBLP. As an aside, because in our computational study $p = 3$, we have 313 features in total. Some of these features rely on the characteristics of the objective functions such as the magnitude and the number of positive, negative and zero coefficients. We also consider features that establish a relationship between the objective functions using some normalization techniques, e.g., the pre-ordering approach used to order the objective functions (see Section 4). Other features are created based on some mathematical and statistical computations that link the objective functions with the technological coefficients and the right-hand-side values.

We also define features based on the area of the projected criterion space i.e., the corresponding $(p - 1)$ -dimensional criterion space, that needs to be explored when one of the objectives is selected for conducting the projection. Note that, to compute such an area, several single-objective binary linear programming problems need to be solved. However, in order to reduce the complexity of the features extraction, we compute an approximation of the area-to-explore by optimizing the linear relaxation of the problems. Additionally, we create features in which the basic characteristics of an instance are described, e.g., size, number of variables, and number of constraints.

The main idea of the features is to generate as much information as possible in a simple way. We accomplish this by computing all the proposed features in polynomial time for a given instance. The features are normalized using a t-statistic score. Normalization is performed by aggregating subsets of features computed from a similar source. Finally, the values of the normalized feature matrix are distributed approximately between -1 and 1. Interested readers can find a detailed explanation of the features in Appendix A.

5.2. Labels

Based on our research goal, i.e., simulating the selection of the best objective, we propose a multi-class integer labeling scheme for each instance, where $y \in \Omega$ is the label of the instance and $\Omega = \{1, 2, \dots, p\}$ is the domain of the labels. The value of y^i classifies the instance i with a label that indicates the index of the best objective function for projection based on the running time of the KSA (when generating the entire nondominated frontier of the instance). The label of each instance is assigned as follows:

$$y^i \in \arg \min_{j \in \{1, \dots, p\}} \{\text{RunningTime}_j^i\}, \quad (2)$$

where RunningTime_j^i is the running time of the instance i when the objective function j is used for projecting the criterion space.

6. Best subset selection of features

It is easy to observe that by introducing more (linearly independent) features and re-training an ML model (to optimality) its prediction error, i.e., $\text{error} = 1 - \text{accuracy}$, on the training set decreases and eventually it becomes zero. This is because in that case we are providing a larger degree of freedom to the ML model. However, this is not necessarily the case for the testing set. In other words, by introducing more features, the ML model that will be obtained is often overfitted to the training set and does not perform well on the testing set. So, the underlying idea of the best subset selection of features is to avoid the issue of overfitting. However, the key point is that in a real-world scenario we do not have access to the testing set. So, selecting the best subset of features should be done based on the information obtained from the training set.

In light of the above, studying the trade-off between the number of features and the prediction error of an ML model on the training set is helpful for selecting the best subset of features (Charkhgard and Eshragh 2019). However, computing such a tradeoff using exact methods is difficult in practice since the total number of subsets (of features) is an exponential function of the number of features. Therefore, in this section, we introduce a bi-objective optimization-based heuristic for selecting the best subset of features. The proposed approach has two phases:

- *Phase I*: In the first phase, the algorithm attempts to approximate the tradeoff. Specifically, the algorithm computes an approximated nondominated frontier of a bi-objective optimization problem in which its conflicting objectives are minimizing the number of features and minimizing the prediction error on the training set.
- *Phase II*: In the second phase, the algorithm selects one of the approximated nondominated points and its corresponding MSVM model to be used for prediction on the testing set.

We first explain Phase I. To compute the approximated nondominated frontier, we run MSVM iteratively on a training set. In each iteration, one approximated nondominated point will be generated. The approximated nondominated point obtained in iteration t is denoted by (k_t, e_t) where k_t is the number of features in the corresponding prediction

model (obtained by MSVM) and e_t is the prediction error of the corresponding model on the training set.

To compute the first nondominated point, the proposed approach/algorithm assumes that all features are available and it runs the MSVM to obtain the parameters of the prediction model. We denote by W^t the parameter of the prediction model obtained by MSVM in iteration t . Note that W^t is a $p \times k_t$ matrix where p is the number of objectives. Now consider an arbitrary iteration t . The algorithm will explore the parameters of the prediction model obtained in the previous iteration by MSVM, i.e., W^{t-1} , and will remove the least important feature based on W^{t-1} . Hence, because of removing one feature, we have that $k_t = k_{t-1} - 1$. Specifically, each column of matrix W^{t-1} is associated to a feature. Therefore, the algorithm computes the standard deviation of each column independently. The feature with the minimum standard deviation will be selected and removed in iteration t . Note that MSVM creates a model for each objective function and that is the reason that matrix W^{t-1} has p rows. So, if the standard deviation of a column in the matrix is zero then we know that the corresponding feature is contributing exactly the same in all p models and therefore it can be removed. So, we observe that the standard deviation plays an important role in identifying the least important feature.

Overall, after removing the least important feature, MSVM should be run again for computing W^t and e_t . The algorithm for computing the approximated nondominated frontier terminates as soon as $k^t = 0$. A detailed description of the algorithm for computing the approximated nondominated frontier can be found in Algorithm 1.

In the second phase, we select an approximated nondominated point. However, before doing so, it is worth mentioning that MSVM can take a long time to compute W^t in each iteration of Algorithm 1. So, to avoid this issue, users usually terminate MSVM before it reaches to an optimal solution by imposing some termination conditions including a relative optimality gap tolerance and adjusting the so-called *regularization* parameter (see Crammer and Singer (2001) and Tsochantaridis et al. (2004) for details). In this study, we set the tolerance to 0.1 and the regularization parameter to 5×10^4 (since we numerically observed that MSVM performs better in this case). Such limitations obviously impact the prediction error that will be obtained on the training set, i.e., e^t . So, some of the points that will be reported by Algorithm 1 may dominate each other. Therefore, in Phase II, we

Algorithm 1: Phase I: Computing an approximated nondominated frontier

input: Training set, The set of features

```
1 Queue.create(Q)
2  $t \leftarrow 1$ 
3  $k_t \leftarrow$  The initial number of features
4 while  $k_t \neq 0$  do
5   if  $t \geq 1$  then
6     Find the least important feature from  $W^{t-1}$  and remove it from the set of features
7      $k_t \leftarrow k_{t-1} - 1$ 
8   Compute  $W^t$  by applying MSVM on the training set using the current set of features
9   Compute  $e_t$  by applying the obtained prediction model associated with  $W^t$  on the training set
10  Q.add(( $k_t, e_t$ ))
11   $t \leftarrow t + 1$ 
12 return Q
```

first remove the dominated points. In the remaining of this section we assume that there is no dominated point in the approximated nondominated frontier.

Next, the proposed approach selects an approximated nondominated point that has the minimum Euclidean distance from the (imaginary) *ideal* point, i.e., an imaginary point in the criterion space that has the minimum number of features as well as the minimum prediction error. Such a technique is a special case of optimization over the frontier (Abbas and Chaabane 2006, Jorge 2009, Boland et al. 2017a, Sierra-Altamiranda and Charkhgard 2019). We note that in bi-objective optimization, the ideal point can be computed easily based on the endpoints of the (approximated) nondominated frontier. Let (k', e') and (k'', e'') be the two endpoints in which $k' < k''$ and $e' > e''$. In this case, the ideal point is (k', e'') . Note too that because the first and second objectives have different scales, in this study, we first normalize all approximated nondominated points before selecting a point. Let (k, e) be an arbitrary approximated nondominated point. After normalization, this point will be as follows:

$$\left(\frac{k - k'}{k'' - k'}, \frac{e - e''}{e' - e''} \right).$$

Observe that the proposed normalization technique ensures that the value of each component of a point will be between 0 and 1. As a consequence, in this case, the nominalized ideal point will be always $(0, 0)$. We will discuss about the effectiveness of our proposed best subset selection approach in the next section.

7. A computational study

In this section, we conduct an extensive computational study to evaluate the performance of the KSA when the proposed ML technique is used for learning to project. We generate 1000 tri-objective AP instances and also 1000 tri-objective KP instances based on the procedures described by Kirlik and Sayın (2014). Since there are three objectives, we compute the entire representation of the nondominated frontier three times for each instance using the KSA; In each time, a different objective function will be selected for projection. We employ CPLEX 12.7 as the single-objective binary linear programming solver. All computational experiments are carried out on a Dell PowerEdge R630 with two Intel Xeon E5-2650 2.2 GHz 12-Core Processors (30MB), 128GB RAM, and the RedHat Enterprise Linux 6.8 operating system. We only use a single thread for all experiments.

Our experiments are divided into three parts. In the first part, we run our approach over the entire set of instances using 80% of the data as the training set and 20% of the data as the testing set. The second part evaluates the prediction models obtained in the first part on a reduced testing set. In other words, the training set is as same as the one in the first part but the testing set is smaller. Specifically, if it does not really matter which objective function to be selected for projection (in terms of solution time) for an instance in the testing set of the first part then we remove it. Obviously one can think of such instances as tie cases. In the third part of our experiments, we extend the concept of reduced testing set to the training set. That is, we remove not only the tie cases from the testing set but also from the training set. In general, the goal of reducing testing set and/or training set is to improve the overall accuracy of the prediction model.

At the end of the computational study, we replace MSVM by Random Forest in the proposed ML framework to show the performance of another learning technique. We note that in this computational study, we do not report any time for our proposed ML framework because the aggregated time of generating the features, learning process, and predictions for all 1000 instances of a class of optimization problem, i.e., AP and KP, is around 50 seconds. This implies that on average almost 0.05 seconds are spent on each instance.

7.1. Complete training and testing sets

The first part of our experiments are done on the entire training and testing sets. For each class of optimization problems, i.e., KP and AP, the proposed subset selection approach is run on its corresponding training set. The proposed approach obtains the best subset of

features and its corresponding prediction model for each class of instances. Before providing detailed explanations about the accuracy of such a prediction model on the testing set, it is necessary to show that the proposed bi-objective optimization approach for selecting the best subset of features is indeed effective.

Figure 1 shows the approximated nondominated frontier (obtained during the course of the proposed best subset selection approach) for each class of optimization problems. In this figure, small (red) plus symbols are the outputs of Algorithm 1. The (black) square on the vertical axis shows the ideal point and finally the (yellow) square on the approximated nondominated frontier shows the selected point by the proposed method. First note that we introduced 313 generic features in this paper but the tail of the approximated nondominated frontier in Figure 1 clearly shows that not all 313 features are used. This is because some of the 313 features are not applicable to all classes and will be removed automatically before running the proposed best subset selection approach.

We observe from Figure 1 that, overall, by introducing more features the prediction error decreases for the training set. It is evident when all features are included the accuracy, i.e., $1 - error$, for the training set is around 59.5% and 70% for AP and KP instances, respectively. Of course this is not surprising because the learning procedure will be done based on the training set and introducing more features gives a larger degree of freedom to the learning model. However, this is not necessarily the case for the testing set. Basically, by introducing more features, we may raise the issue of *overfitting*, i.e., the prediction error is small for the training set but large for the testing set.

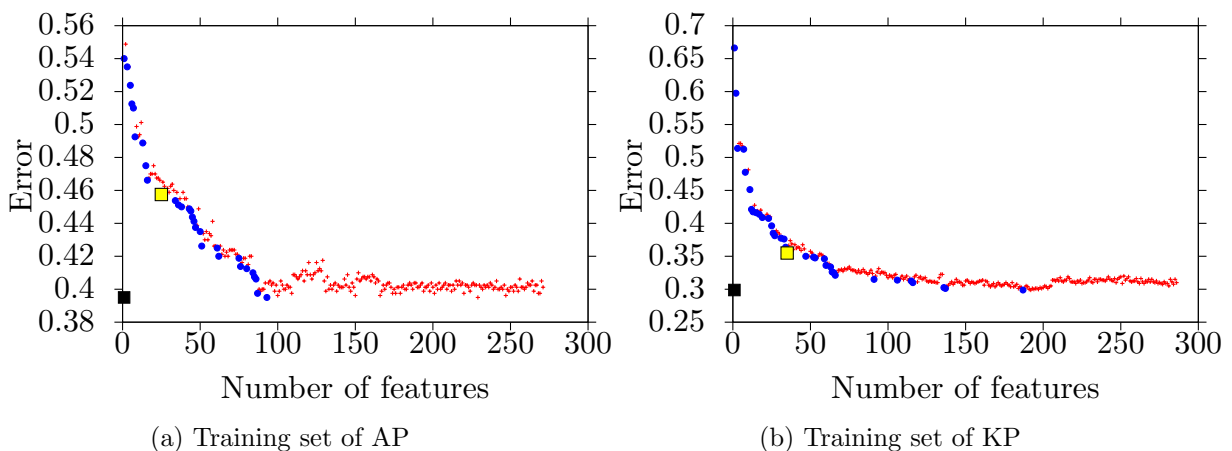


Figure 1 An illustration of the performance of the proposed approach for selecting the best subset of features on the complete training set

To show this, for each of the points (other than the ideal point) in Figure 1, we have plotted its corresponding point for the testing set in Figure 2. Specifically, for each point in Figure 1, we run its corresponding model on the testing set to compute the error. From Figure 2 we observe that the error highly fluctuates. In fact, it is evident that for AP instances, the prediction model that has around 40 features is the best prediction model. Similarly, from Figure 2, it is evident that for KP instances, the prediction model that has around 25 features is the best prediction model.

Note that in practice, we do not have access to the testing set. So, we should select the best subset of features only based on the training set. Therefore, the goal of any best subset selection technique is to identify the prediction model that is (almost) optimal for the testing set based on the existing data set, i.e., the training set. From Figure 2, we observe that our proposed best subset selection heuristic has such a desirable characteristic in practice. We see that the selected model, i.e., the (yellow) square, is nearly optimal. In fact the proposed approach has selected a prediction model with the accuracy of around 50% and 55% for AP and KP instances, respectively. This implies that the absolute difference between the accuracy of the model selected by the proposed subset selection approach and the accuracy of the optimal model is almost 3% and 5% for AP and KP instances, respectively. We note that for both classes of instances less than 50 features exist in the model selected by the proposed approach. Overall, these results are promising due to the fact that the (expected) probability of randomly picking the correct objective function to project is $\frac{1}{p}$, i.e., around 33.3% for the tri-objective instances.

Table 2 Accuracy and average time decrease of testing set when using the proposed ML technique (for the case of the complete training and testing sets)

Type	Vars	Accuracy	Time Decrease		<u>ML vs. Rand</u> <u>Best vs. Rand</u>
			ML vs. Rand	Best vs. Rand	
AP	20 × 20	55.56%	1.29%	2.33%	55.43%
	25 × 25	44.74%	0.67%	1.65%	40.57%
	30 × 30	41.30%	0.18%	1.48%	12.20%
	35 × 35	52.08%	0.44%	1.69%	25.81%
	40 × 40	56.25%	1.07%	1.65%	64.99%
Avg		49.50%	0.68%	1.74%	38.90%
KP	60	63.64%	9.12%	11.03%	82.68%
	70	45.65%	2.01%	10.14%	19.79%
	80	56.10%	4.59%	11.42%	40.22%
	90	58.82%	8.05%	13.27%	60.68%
	100	51.43%	5.09%	10.34%	49.24%
Avg		55.00%	5.67%	11.17%	50.77%




Figure 2 An illustration of the performance of the proposed approach for the best subset selection of features on the complete testing set

We now discuss about the performance of the selected model in detail for each class of optimization problems. Table 2 summarizes our findings. In this table, the column labeled ‘Accuracy’ shows the average percentage of the prediction accuracy of the selected model for different subclasses of instances. Note that as mentioned in Introduction, each subclass has 200 instances. The column labeled ‘ML vs. Rand’ shows the average percentage of decrease in solution time when ML technique is used compared to randomly picking an objective function for projecting. The column labeled ‘Best vs. Rand’ shows the average percentage of decrease in solution time when the best objective function is selected for projection compared to randomly picking an objective function for projecting. Finally, column labeled ‘ $\frac{\text{ML vs. Rand}}{\text{Best vs. Rand}}$ ’ shows the percentage of ‘ML vs. Rand’ to ‘Best vs. Rand’.

Overall, we observe that our ML method improves the computational time in all testing sets. For AP instances, the improvement is around 0.68% on average which is small. However, we should note that in the ideal case, we could obtain around 1.74% improvement in time on average for such instances. So, the improvement obtained with the proposed ML technique is 38.9% of the ideal case. For largest subclass of AP instances, this number is around 64.99%. For the KP instances, the results are even more promising since the amount of improvement in solution time is around 5.67% on average. In the ideal case, we could obtain an average improvement of 11.17% for such instances. So, the improvement obtained with the proposed ML technique is 50.77% of the ideal case.

7.2. Complete training set and reduced testing set

In this section, we test the performance of the model obtained in Section 7.1 on a reduced testing set. Specifically, we remove the instances that can be considered as tie cases, i.e., those in which the solution time does not change significantly (relative to other instances) when different objective functions are selected for projection. To reduce the testing set we apply the following steps:

- *Step 1:* We compute the *time range* of each instance, i.e., the difference between the best and worst solution times that can be obtained for the instance when different objective functions are considered for projection.
- *Step 2:* For each subclass of instances, i.e., those with the same number of decision variables, we compute the standard deviation and the minimum of time ranges in that subclass.
- *Step 3:* We eliminate an instance, i.e., consider it as a tie case, if its time range is not greater than the sum of the minimum and standard deviation of time ranges in its associated subclass.

As a result of the procedure explained above, the testing set was reduced by 35.5% for AP instances and by 17.5% for KP instances. Table 3 summarizes our findings for the reduced testing set.

Table 3 Accuracy and average time decrease of testing set when using the proposed ML technique (for the case of the complete training set and the reduced testing set)

Type	Vars	Accuracy	Time Decrease		ML vs. Rand Best vs. Rand
			ML vs. Rand	Best vs. Rand	
AP	20 × 20	58.06%	1.48%	2.54%	58.10%
	25 × 25	60.00%	0.92%	2.23%	41.32%
	30 × 30	40.74%	0.12%	1.96%	6.02%
	35 × 35	57.89%	0.54%	1.97%	27.17%
	40 × 40	69.57%	1.64%	2.15%	76.09%
Avg		56.59%	0.90%	2.16%	41.73%
KP	60	71.05%	10.11%	11.96%	84.52%
	70	45.95%	2.03%	11.48%	17.67%
	80	57.58%	5.30%	13.48%	39.31%
	90	62.96%	10.21%	15.57%	65.59%
	100	60.00%	6.32%	11.33%	55.77%
Avg		59.39%	6.66%	12.63%	52.74%

Observe that the accuracy of the prediction models has increased significantly for the reduced testing set. Specifically, it has reached to around 56.59% and 59.39% on average for AP and KP instances, respectively. Since the eliminated instances are considered as tie

cases, we can assume that they are also success cases for the prediction model. So, by considering such success cases, the prediction accuracy will increase to $56.59 \times (1 - 0.355) + 35.5 \cong 72\%$ and $59.39 \times (1 - 0.175) + 17.5 \cong 66.5\%$ for AP and KP instances, respectively. In terms of computational time, we also observe (from Table 3) an improvement of around 0.90% and 6.66% on average for AP and KP instances, respectively. This amount of improvement is about 41.73% and 52.74% of the ideal scenarios (on average) for AP and KP instances, respectively.


7.3. Reduced training and testing sets

Due to promising results obtained in Section 7.2, it is natural to ask whether we can see even more improvement if we reduce not only the testing set but also the training set. Therefore, in this section, we eliminate the tie cases using the same procedure discussed in Section 7.2 from both training and testing sets. By doing so, the size of the training+testing set was reduced by 37% and 18% for AP and KP instances, respectively.

It is evident that due to the change in the training set, we need to apply our proposed approach for best subset selection of features again. So, similar to Section 7.1, Figure 3 shows the approximated nondominated frontier for each class of optimization problems based on the reduced training set. By comparing the ideal points in Figures 1 and 3, an immediate improvement in the (ideal) accuracy can be observed. In fact the absolute difference between the error of the ideal points (in these figures) is around 12% and 7% for AP and KP instances, respectively. Similar improvements can be observed by comparing the selected approximated nondominated points in Figures 1 and 3.

Similar to Section 7.1, for each of the points (other than the ideal point) in Figure 3, we have plotted its corresponding point for the testing set in Figure 4. We again observe that the selected model, i.e., the (yellow) square, is nearly optimal for both classes of optimization problems. In fact the proposed approach has selected a prediction model with the accuracy of around 52% and 62% for AP and KP instances, respectively. This implies that the absolute difference between the accuracy of the model selected by the proposed approach and the accuracy of the optimal model is almost 5% and 3% for AP and KP instances, respectively.


A summary of the results of this last experiment can be found in Table 4. Observe that the average prediction accuracy on the testing set for the experimental setting in this section, i.e., reduced training and testing sets, has improved significantly for KP instances



(a) Training set of AP

(b) Training set of KP

Figure 3 An illustration of the performance of the proposed approach for selecting the best subset of features on the reduced training set



(a) Testing set of AP

(b) Testing set of KP

Figure 4 An illustration of the performance of the proposed approach for the best subset selection of features on the reduced testing set

compared to the results given in Sections 7.1 and 7.2. However for the instances of AP problem, the average prediction accuracy in this section is only better than the one presented in Section 7.1. Overall, the average prediction accuracy is 52.38% and 61.59% for AP and KP instances when using reduced training and testing sets. By considering the tie cases as success events, the projected accuracy increases up to 70% and 68.5% for AP and KP instances, respectively. The importance of such an increase in the accuracy is highlighted by the time decrease percentages given in Table 4, which is over 1% for AP instances and is near 8% for KP instances. In fact, for the largest subclass of AP instances, the average time improvement of 1.6% is equivalent to almost 110 seconds on average. Similarly, for

Table 4 Accuracy and average time decrease of testing set when using the proposed ML technique (for the case of the reduced training and testing sets)

Type	Vars	Accuracy	Time Decrease		<u>ML vs. Rand</u> <u>Best vs. Rand</u>
			ML vs. Rand	Best vs. Rand	
AP	20 × 20	58.33%	0.82%	2.14%	38.41%
	25 × 25	52.38%	1.17%	2.39%	48.76%
	30 × 30	55.17%	0.67%	1.93%	34.46%
	35 × 35	45.45%	1.10%	2.53%	43.35%
	40 × 40	52.63%	1.60%	2.87%	55.61%
Avg		52.38%	1.03%	2.35%	43.99%
KP	60	60.61%	6.75%	12.16%	55.51%
	70	55.56%	6.79%	12.03%	56.43%
	80	60.00%	12.14%	16.91%	71.79%
	90	72.73%	9.22%	12.26%	75.17%
	100	62.50%	4.70%	10.56%	44.50%
Avg		61.59%	7.61%	12.64%	60.18%

the largest subclass of KP instances, the time improvement of 4.7% is around 490 seconds on average.

7.4. Replacing MSVM by Random Forest

One main reason that we used MSVM in this study is that (as shown in the previous sections) it performs well in practice for the purpose of this study. However, another critical reason is the fact that MSVM creates a matrix of parameters denoted by W^t in each iteration. This matrix has p rows where p is the number of objective functions. In other words, for each objective function, MSVM creates a specific model for predicting which one should be used for projection in KSA. This characteristic is desirable because it allowed us to develop a custom-built bi-objective heuristic for selecting the best subset of features. Specifically, as discussed in Section 6, this characteristic is essential for identifying the least important feature in each iteration of Algorithm 1. However, applying such a procedure on other ML techniques is not trivial.

In light of the above, in this section we replace MSVM by Random Forest within the proposed machine learning framework. However, we simply use the best subset of features selected by MSVM and then feed it to Random Forest for training and predicting. To implement Random Forest we use `scikit-learn` library in Python (Pedregosa et al. 2011). Table 5 shows a comparison between the prediction accuracy of MSVM and Random Forest under three experimental settings described in Sections 7.1-7.3. In other words, Setting 1 corresponds to the complete training and testing sets; Setting 2 corresponds to the complete training set and reduced testing sets; Finally, Setting 3 refers to the reduced training and testing sets.

In this table, columns labeled ‘Increase’ show the average percentage of increase in the prediction accuracy of the Random Forest compared to MSVM. Observe from these columns that the reported numbers are mostly negative. This implies that, in general, MSVM outperforms Random Forest in terms of prediction accuracy. For example, in Setting 3, we observe that the accuracy of Random Forest is around 18.63% and 22.88% worse than the accuracy of MSVM on average. This experiment clearly shows the advantage of using MSVM in the proposed ML framework.

Table 5 A performance comparison between MSVM and Random Forest on a testing set

Type	Vars	Setting 1		Setting 2		Setting 3	
		Accuracy	Increase	Accuracy	Increase	Accuracy	Increase
AP	20 × 20	61.11%	9.99%	64.52%	11.12%	36.36%	-37.66%
	25 × 25	42.11%	-5.89%	50.00%	-16.67%	26.67%	-49.09%
	30 × 30	45.65%	10.54%	37.04%	-9.09%	52.63%	-4.60%
	35 × 35	62.50%	20.01%	65.79%	13.65%	70.59%	55.31%
	40 × 40	50.00%	-11.11%	73.91%	6.24%	50.00%	-5.00%
Avg		52.50%	6.06%	59.69%	5.48%	42.62%	-18.63%
KP	60	47.73%	-25.00%	47.37%	-33.33%	41.18%	-32.06%
	70	50.00%	9.53%	56.76%	23.52%	50.00%	-10.01%
	80	48.78%	-13.06%	54.55%	-5.27%	61.90%	3.17%
	90	44.12%	-25.00%	55.56%	-11.76%	47.06%	-35.30%
	100	37.14%	-27.78%	43.33%	-27.78%	45.95%	-26.49%
Avg		46.00%	-16.36%	51.52%	-13.26%	47.50%	-22.88%

8. Conclusions and future research

We presented a multi-class support vector machine based approach to enhance exact multi-objective binary linear programming algorithms. Our approach simulates the best selection of objective function to be used for projection in the KSA in order to improve its computational time. We introduced a pre-ordering approach for the objective functions in the input file for the purpose of standardizing the vector of features. Moreover, we introduced a bi-objective optimization approach for selecting the best subset of features in order to overcome overfitting. By conducting an extensive computational , we showed that reaching to the prediction accuracy of around 70% is possible for instances of tri-objective AP and KP. It was shown that such a prediction accuracy results in a decrease of over 12% in the computational time for some instances.

Overall, we hope that the simplicity of our proposed ML technique and its promising results encourage more researchers to use ML techniques for improving multi-objective

optimization solvers. Note that, in this paper, we studied the problem of learning to project in a static setting, i.e., before solving an instance we predict the best objective function and use it during the course of the search. So, one future research direction of this study would be finding a way to employ the proposed learning-to-project technique in a dynamic setting, i.e., at each iteration in the search process we predict the best projected objective and use it. Evidently, this may result in developing new algorithms that have not yet been studied in the literature of multi-objective optimization.

References

- Abbas M, Chaabane D (2006) Optimizing a linear function over an integer efficient set. *European Journal of Operational Research* 174(2):1140–1161.
- Alvarez AM, Louveaux Q, Wehenkel L (2017) A machine learning-based approximation of strong branching. *INFORMS Journal on Computing* 29(1):185–195.
- Bertsimas D, King A, Mazumder R, et al. (2016) Best subset selection via a modern optimization lens. *The annals of statistics* 44(2):813–852.
- Boland N, Charkhgard H, Savelsbergh M (2015a) A criterion space search algorithm for biobjective integer programming: The balanced box method. *INFORMS Journal on Computing* 27(4):735–754.
- Boland N, Charkhgard H, Savelsbergh M (2015b) A criterion space search algorithm for biobjective mixed integer programming: The triangle splitting method. *INFORMS Journal on Computing* 27(4):597–618.
- Boland N, Charkhgard H, Savelsbergh M (2016) The L-shape search method for triobjective integer programming. *Mathematical Programming Computation* 8(2):217–251.
- Boland N, Charkhgard H, Savelsbergh M (2017a) A new method for optimizing a linear function over the efficient set of a multiobjective integer program. *European Journal of Operational Research* 260(3):904 – 919.
- Boland N, Charkhgard H, Savelsbergh M (2017b) The quadrant shrinking method: A simple and efficient algorithm for solving tri-objective integer programs. *European Journal of Operational Research* 260(3):873 – 885.
- Bottou L (2010) Large-scale machine learning with stochastic gradient descent. *Proceedings of COMP-STAT'2010*, 177–186 (Springer).
- Breiman L (2001) Random forests. *Machine learning* 45(1):5–32.
- Charkhgard H, Eshragh A (2019) A new approach to select the best subset of predictors in linear regression modeling: bi-objective mixed integer linear programming. *ANZIAM journal*. Available online. <https://doi.org/10.1017/S1446181118000275>.
- Crammer K, Singer Y (2001) On the algorithmic implementation of multiclass kernel-based vector machines. *Journal of machine learning research* 2(Dec):265–292.

- Dächert K, Gorski J, Klamroth K (2012) An augmented weighted Tchebycheff method with adaptively chosen parameters for discrete bicriteria optimization problems. *Computers & Operations Research* 39:2929–2943.
- Dächert K, Klamroth K (2015) A linear bound on the number of scalarizations needed to solve discrete tricriteria optimization problems. *Journal of Global Optimization* 61(4):643–676.
- Fattahi A, Turkay M (2017) A one direction search method to find the exact nondominated frontier of biobjective mixed-binary linear programming problems. *European Journal of Operational Research* . Available online.
- He H, Daume III H, Eisner JM (2014) Learning to search in branch and bound algorithms. *Advances in neural information processing systems*, 3293–3301.
- Jorge JM (2009) An algorithm for optimizing a linear function over an integer efficient set. *European Journal of Operational Research* 195(1):98–103.
- Khalil EB, Dilkina B, Nemhauser GL, Ahmed S, Shao Y (2017) Learning to run heuristics in tree search. *Proceedings of the international joint conference on artificial intelligence. AAAI Press, Melbourne, Australia*.
- Khalil EB, Le Bodic P, Song L, Nemhauser GL, Dilkina BN (2016) Learning to branch in mixed integer programming. *AAAI*, 724–731.
- Kirlik G, Sayın S (2014) A new algorithm for generating all nondominated solutions of multiobjective discrete optimization problems. *European Journal of Operational Research* 232(3):479 – 488.
- Le QV, Ngiam J, Coates A, Lahiri A, Prochnow B, Ng AY (2011) On optimization methods for deep learning. *Proceedings of the 28th International Conference on International Conference on Machine Learning*, 265–272 (Omnipress).
- Lokman B, Köksalan M (2013) Finding all nondominated points of multi-objective integer programs. *Journal of Global Optimization* 57(2):347–365.
- Özlen M, Azizoğlu M (2009) Multi-objective integer programming: A general approach for generating all non-dominated solutions. *European Journal of Operational Research* 199:25–35.
- Özlen M, Burton BA, MacRae CAG (2013) Multi-objective integer programming: An improved recursive algorithm. *Journal of Optimization Theory and Applications* . 10.1007/s10957-013-0364-y.
- Özpeynirci Ö, Köksalan M (2010) An exact algorithm for finding extreme supported nondominated points of multiobjective mixed integer programs. *Management Science* 56(12):2302–2315.
- Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay E (2011) Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12:2825–2830.
- Prinzie A, Van den Poel D (2008) Random forests for multiclass classification: Random multinomial logit. *Expert systems with Applications* 34(3):1721–1732.

- Przybylski A, Gandibleux X (2017) Multi-objective branch and bound. *European Journal of Operational Research* 260(3):856 – 872.
- Przybylski A, Gandibleux X, Ehrgott M (2010) A two phase method for multi-objective integer programming and its application to the assignment problem with three objectives. *Discrete Optimization* 7(3):149 – 165.
- Roth D, Yih Wt (2005) Integer linear programming inference for conditional random fields. *Proceedings of the 22nd international conference on Machine learning*, 736–743 (ACM).
- Sabharwal A, Samulowitz H, Reddy C (2012) Guiding combinatorial optimization with uct. *International Conference on Integration of Artificial Intelligence (AI) and Operations Research (OR) Techniques in Constraint Programming*, 356–361 (Springer).
- Sierra-Altamiranda A, Charkhgard H (2019) A new exact algorithm to optimize a linear function over the set of efficient solutions for bi-objective mixed integer linear programming. *INFORMS Journal on Computing* . To appear.
- Snoek J, Larochelle H, Adams RP (2012) Practical bayesian optimization of machine learning algorithms. *Advances in neural information processing systems*, 2951–2959.
- Soylu B, Yıldız GB (2016) An exact algorithm for biobjective mixed integer linear programming problems. *Computers & Operations Research* 72:204–213.
- Sra S, Nowozin S, Wright SJ (2012) *Optimization for machine learning* (MIT Press).
- Tibshirani R (1996) Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B* 58:267–288.
- Tsochantaridis I, Hofmann T, Joachims T, Altun Y (2004) Support vector machine learning for interdependent and structured output spaces. *Proceedings of the twenty-first international conference on Machine learning*, 104 (ACM).
- Vincent T, Seipp F, Ruzika S, Przybylski A, Gandibleux X (2013) Multiple objective branch and bound for mixed 0-1 linear programming: Corrections and improvements for biobjective case. *Computers & Operations Research* 40(1):498–509.

Appendix A: List of Features

In this section, we present all $5p^2 + 106p - 50$ features that we used in this study. Note that since $p = 3$ in our computational study, the total number of features are 313 features. For convenience, we partition all the proposed features into some subsets and present them next. We use the letter F to represent each subset.

The first subset of features is $F^1 = \{\mathbf{c}_1^\top \tilde{\mathbf{x}}, \mathbf{c}_2^\top \tilde{\mathbf{x}}, \dots, \mathbf{c}_p^\top \tilde{\mathbf{x}}\}$, which will be automatically computed during the pre-ordering process. To incorporate the impact of the size of an instance in the learning process, we introduce $F^2 = \{n\}$, $F^3 = \{m\}$, and $F^4 = \{\text{density}(A)\}$. To incorporate the impact of zero coefficients of the objective functions in the learning process, for each $i \in \{1, \dots, p\}$ we introduce:

$$F_i^5 = \{\text{size}(S_i^5)\},$$

where $S_i^5 = \{c \in \mathbf{c}_i : c = 0\}$. To incorporate the impact of positive coefficients of the objective functions in the learning process, for each $i \in \{1, \dots, p\}$ we introduce:

$$F_i^6 = \{\text{size}(S_i^6), \text{Avg}(S_i^6), \text{Min}(S_i^6), \text{Max}(S_i^6), \text{Std}(S_i^6), \text{Median}(S_i^6)\},$$

where $S_i^6 = \{c \in \mathbf{c}_i : c > 0\}$. To incorporate the impact of negative coefficients of the objective functions in the learning process, for each $i \in \{1, \dots, p\}$ we introduce:

$$F_i^7 = \{\text{size}(S_i^7), \text{Avg}(S_i^7), \text{Min}(S_i^7), \text{Max}(S_i^7), \text{Std}(S_i^7), \text{Median}(S_i^7)\},$$

where $S_i^7 = \{c \in \mathbf{c}_i : c < 0\}$. To establish a relation between the objective functions and A in the learning process, for each $i \in \{1, \dots, p\}$ we introduce:

$$F_i^8 = \{\text{Avg}(S_i^8), \text{Min}(S_i^8), \text{Max}(S_i^8), \text{Std}(S_i^8), \text{Median}(S_i^8)\},$$

where $S_i^8 = \cup_{j \in \{1, \dots, m\}} \{\mathbf{c}_i^\top \mathbf{a}_j^\top\}$ and \mathbf{a}_j represents row j of matrix A . For the same purpose, for each $i \in \{1, \dots, p\}$ we also introduce:

$$F_i^9 = \mathbf{c}_i^\top \times A^\top \times \mathbf{b}.$$

For each $j \in \{1, \dots, m\}$, let $b'_j := b_j + 1$ if $b_j \geq 0$ and $b'_j := b_j - 1$ otherwise. To establish a relation between the positive and negative coefficients in the objective functions and \mathbf{b} in the learning process, for each $i \in \{1, \dots, p\}$ and $k \in \{10, 11\}$ we introduce:

$$F_i^k = \{\text{Avg}(S_i^k), \text{Min}(S_i^k), \text{Max}(S_i^k), \text{Std}(S_i^k), \text{Median}(S_i^k)\},$$

where $S_i^{10} = \cup_{j \in \{1, \dots, m\}} \{\frac{\sum_{c \in S_i^6} c}{b'_j}\}$ and $S_i^{11} = \cup_{j \in \{1, \dots, m\}} \{\frac{\sum_{c \in S_i^7} c}{b'_j}\}$.

For each $i \in \{1, \dots, p\}$, let $l_i := \min_{\mathbf{x} \in \mathcal{X}_{LR}} z_i(\mathbf{x})$ and $u_i := \max_{\mathbf{x} \in \mathcal{X}_{LR}} z_i(\mathbf{x})$ where \mathcal{X}_{LR} is the linear programming relaxation of \mathcal{X} . To incorporate the impact of the volume of the search region in a projected criterion space, i.e., a $(p-1)$ -dimensional criterion space, in the learning process, for each $i \in \{1, \dots, p\}$ we introduce:

$$F_i^{12} = \prod_{j \in \{1, \dots, p\} \setminus \{i\}} (u_j - l_j).$$

Let $\bar{c}'_i := \frac{\sum_{c \in \mathbf{c}_i} |c|}{n}$ be the average of the absolute values of the elements in the objective $i \in \{1, \dots, p\}$. Note that we understand that \mathbf{c}_i is a vector (and not a set) and hence $c \in \mathbf{c}_i$ is not a well-defined mathematical notation. However, for simplicity, we keep this notation as is and basically treat each component of the vector as an element.

We also introduce some features that measures the size of an instance in an indirect way. Specifically, for each $i \in \{1, \dots, p\}$ and $k \in \{13, 14, 15\}$ we introduce:

$$F_i^k = \{\text{Avg}(S_i^k), \text{Min}(S_i^k), \text{Max}(S_i^k), \text{Std}(S_i^k), \text{Median}(S_i^k)\},$$

to measure n , p and m , respectively, where

$$S_i^{13} := \cup_{j \in \{1, \dots, p\} \setminus \{i\}} \left\{ \frac{\sum_{c \in \mathbf{c}_i} |c|}{\bar{c}'_j + 1} \right\},$$

$$S_i^{14} := \cup_{k \in \{1, \dots, n\}} \left\{ \frac{\sum_{j \in \{1, \dots, p\} \setminus \{i\}} |c_{jk}|}{|c_{ik}| + 1} \right\},$$

and

$$S_i^{15} := \cup_{k \in \{1, \dots, m\}} \left\{ \frac{\sum_{j=1}^m |a_{jk}|}{|c_{ik}| + 1} \right\}.$$

Motivated by the idea using the product of two variables for studying the interaction effect between them, for each $i \in \{1, \dots, p\}$, we introduce:

$$F_i^{16} = \{\text{Avg}(S_i^{16}), \text{Min}(S_i^{16}), \text{Max}(S_i^{16}), \text{Std}(S_i^{16}), \text{Median}(S_i^{16})\},$$

where $S_i^{16} = \cup_{j \in \{1, \dots, p\} \setminus \{i\}} \left\{ \sum_{l=1}^n c_{il} c_{jl} \right\}$. Similarly, we also define a subset of features based on the *leverage score* LS_j of the variable $j \in \{1, \dots, n\}$ in the matrix A . Specifically, for each $i \in \{1, \dots, p\}$, we introduce:

$$F_i^{17} = \sum_{j=1}^n c_{ij} LS_j.$$

where $LS_j := \frac{\|\mathbf{a}_j\|^2}{\sum_{l=1}^n \|\mathbf{a}_l\|^2}$ and \mathbf{a}_j represents column j of matrix A for each $j \in \{1, \dots, n\}$. Let $\text{Avg}(C) := \text{Avg}(\mathbf{c}_1, \dots, \mathbf{c}_p)$, $\text{Std}(C) := \text{Std}(\mathbf{c}_1, \dots, \mathbf{c}_p)$, and

$$O := \{(-\infty, -1), (-1, -0.5), (-0.5, 0), (0, 0.5), (0.5, 1), (1, \infty)\}.$$

For each $i \in \{1, \dots, p\}$, we define

$$F_i^{18} = \cup_{(l,u) \in O} \{\text{car}(\mathbf{c}_i^{l,u})\}$$

where

$$\mathbf{c}_i^{l,u} := \{c \in \mathbf{c}_i : \text{Avg}(C) + l \text{Std}(C) \leq c \leq \text{Avg}(C) + u \text{std}(C)\}.$$

The following observation creates the basis of the remaining features.

OBSERVATION 1. Let $\alpha_i > 0$ where $i \in \{1, \dots, p\}$ and β_k where $k \in \{1, \dots, m\}$ be positive constants. For a MOBLP, its equivalent problem can be constructed as follows:

$$\begin{aligned} \min \quad & \left\{ \sum_{j=1}^n \alpha_1 c_{1j} x_j, \dots, \sum_{j=1}^n \alpha_p c_{pj} x_j \right\} \\ \text{s.t.} \quad & \sum_{j=1}^n \beta_k a_{kj} x_j \leq \beta_k b_k \quad \forall k \in \{1, \dots, m\}, \\ & x_j \in \{0, 1\} \quad \forall j \in \{1, \dots, n\} \end{aligned} \tag{3}$$

Observation 1 is critical because it shows that our ML approach should not be sensitive to positive scaling. So, the remaining features are specifically designed to address this issue. Note that the remaining features are similar to the ones that we have already seen before but they are less sensitive to a positive scaling.

Let $c_i^{max} = \max\{|c_{i1}|, \dots, |c_{in}|\}$ and $\bar{c}_i = (\frac{c_{i1}}{c_i^{max}}, \dots, \frac{c_{in}}{c_i^{max}})$. To incorporate the impact of the relative number of zeros, positive, and negative coefficients of the objective functions in the learning process, for each $i \in \{2, \dots, p\}$, we introduce:

$$F_i^{19} = \left\{ \ln \left(1 + \frac{car(\bar{c}_i)^0}{1 + car(\bar{c}_i)^0} \right) \right\},$$

$$F_i^{20} = \left\{ \ln \left(1 + \frac{car(\bar{c}_i)^+}{1 + car(\bar{c}_i)^+} \right) \right\},$$

and

$$F_i^{21} = \left\{ \ln \left(1 + \frac{car(\bar{c}_i)^-}{1 + car(\bar{c}_i)^-} \right) \right\},$$

where $car(\bar{c}_i)^0$ is the number of elements in \bar{c}_i with zero values. Also, $car(\bar{c}_i)^+$ is the number of elements in \bar{c}_i with positive values. Finally, $car(\bar{c}_i)^-$ is the number of elements in \bar{c}_i with negative values.

The following function is helpful for introducing some other features:

$$g(a) = \begin{cases} a + 1 & \text{if } a \geq 0 \\ a - 1 & \text{otherwise} \end{cases}.$$

For each $l \in \{1, \dots, m\}$, let $a_l^{max} = \max\{|a_{l1}|, \dots, |a_{ln}|\}$, $\bar{a}_l = (\frac{a_{l1}}{a_l^{max}}, \dots, \frac{a_{ln}}{a_l^{max}})$ and $\bar{b}_l = \frac{b_l}{a_l^{max}}$. To incorporate the relative impact of the magnitude of objective function coefficients, and constraints in the learning process, for each $i \in \{2, \dots, p\}$ and $k \in \{22, 23, 24, 25, 26, 27\}$, we introduce:

$$F_i^k = \{Avg(S_i^k), Min(S_i^k), Max(S_i^k), Std(S_i^k), Median(S_i^k)\},$$

where

$$S_i^{22} = \left\{ \frac{\bar{c}_{11}}{g(\bar{c}_{i1})}, \dots, \frac{\bar{c}_{1n}}{g(\bar{c}_{in})} \right\},$$

$$S_i^{23} = \left\{ \frac{\bar{c}_{11}^2}{g(\bar{c}_{i1}^2)}, \dots, \frac{\bar{c}_{1n}^2}{g(\bar{c}_{in}^2)} \right\},$$

$$S_i^{24} = \left\{ \sum_{j=1}^n \frac{\bar{c}_{1j} \bar{a}_{1j}}{ng(\bar{c}_{ij})}, \dots, \sum_{j=1}^n \frac{\bar{c}_{1j} \bar{a}_{mj}}{ng(\bar{c}_{ij})} \right\},$$

$$S_i^{25} = \left\{ \sum_{j=1}^n \frac{\bar{c}_{1j} \bar{a}_{1j}}{ng(\bar{c}_{ij})} - \bar{b}_1, \dots, \sum_{j=1}^n \frac{\bar{c}_{1j} \bar{a}_{mj}}{ng(\bar{c}_{ij})} - \bar{b}_m \right\},$$

$$S_i^{26} = \left\{ \sum_{j=1}^n \frac{\bar{c}_{1j}^2 \bar{a}_{1j}}{ng(\bar{c}_{ij}^2)}, \dots, \sum_{j=1}^n \frac{\bar{c}_{1j}^2 \bar{a}_{mj}}{ng(\bar{c}_{ij}^2)} \right\},$$

and

$$S_i^{27} = \left\{ \sum_{j=1}^n \frac{\bar{c}_{1j}^2 \bar{a}_{1j}}{ng(\bar{c}_{ij}^2)} - \bar{b}_1, \dots, \sum_{j=1}^n \frac{\bar{c}_{1j}^2 \bar{a}_{mj}}{ng(\bar{c}_{ij}^2)} - \bar{b}_m \right\}.$$

For the same reason, for each $i \in \{1, \dots, p\}$ and $l \in \{1, \dots, p\} \setminus \{i\}$, we introduce:

$$F_{il}^{28} = \{Avg(S_{il}^{28}), Min(S_{il}^{28}), Max(S_{il}^{28}), Std(S_{il}^{28}), Median(S_{il}^{28})\},$$

where

$$S_{il}^{28} = \left\{ \frac{\bar{c}_{i1}}{g(\bar{c}_{l1})}, \dots, \frac{\bar{c}_{in}}{g(\bar{c}_{ln})} \right\}.$$

Finally, let $\bar{A}_j = \{\bar{a}_{1j}, \dots, \bar{a}_{mj}\}$ for each $j \in \{1, \dots, n\}$. For each $i \in \{2, \dots, p\}$, the following subsets of features are also defined for linking the constraints and objective functions:

$$F_i^{29} = \left\{ \sum_{j=1}^n \frac{\bar{c}_{1j} \text{Avg}(\bar{A}_j)}{ng(\bar{c}_{ij})}, \sum_{j=1}^n \frac{\bar{c}_{1j} \text{Min}(\bar{A}_j)}{ng(\bar{c}_{ij})}, \sum_{j=1}^n \frac{\bar{c}_{1j} \text{Max}(\bar{A}_j)}{ng(\bar{c}_{ij})}, \sum_{j=1}^n \frac{\bar{c}_{1j} \text{Std}(\bar{A}_j)}{ng(\bar{c}_{ij})}, \sum_{j=1}^n \frac{\bar{c}_{1j} \text{Median}(\bar{A}_j)}{ng(\bar{c}_{ij})} \right\},$$

$$F_i^{30} = \left\{ \sum_{j=1}^n \frac{\bar{c}_{1j} \text{Avg}(\bar{A}_j)}{ng(\bar{c}_{ij})} - \text{Avg}(\bar{b}_j), \sum_{j=1}^n \frac{\bar{c}_{1j} \text{Min}(\bar{A}_j)}{ng(\bar{c}_{ij})} - \text{Min}(\bar{b}_j), \sum_{j=1}^n \frac{\bar{c}_{1j} \text{Max}(\bar{A}_j)}{ng(\bar{c}_{ij})} - \text{Max}(\bar{b}_j), \right. \\ \left. \sum_{j=1}^n \frac{\bar{c}_{1j} \text{Std}(\bar{A}_j)}{ng(\bar{c}_{ij})} - \text{Std}(\bar{b}_j), \sum_{j=1}^n \frac{\bar{c}_{1j} \text{Median}(\bar{A}_j)}{ng(\bar{c}_{ij})} - \text{Median}(\bar{b}_j) \right\},$$

$$F_i^{31} = \left\{ \sum_{j=1}^n \frac{\bar{c}_{1j}^2 \text{Avg}(\bar{A}_j)}{ng(\bar{c}_{ij}^2)}, \sum_{j=1}^n \frac{\bar{c}_{1j}^2 \text{Min}(\bar{A}_j)}{ng(\bar{c}_{ij}^2)}, \sum_{j=1}^n \frac{\bar{c}_{1j}^2 \text{Max}(\bar{A}_j)}{ng(\bar{c}_{ij}^2)}, \sum_{j=1}^n \frac{\bar{c}_{1j}^2 \text{Std}(\bar{A}_j)}{ng(\bar{c}_{ij}^2)}, \sum_{j=1}^n \frac{\bar{c}_{1j}^{\text{@}} \text{Median}(\bar{A}_j)}{ng(\bar{c}_{ij}^2)} \right\},$$

$$F_i^{32} = \left\{ \sum_{j=1}^n \frac{\bar{c}_{1j}^2 \text{Avg}(\bar{A}_j)}{ng(\bar{c}_{ij}^2)} - \text{Avg}(\bar{b}_j), \sum_{j=1}^n \frac{\bar{c}_{1j}^2 \text{Min}(\bar{A}_j)}{ng(\bar{c}_{ij}^2)} - \text{Min}(\bar{b}_j), \sum_{j=1}^n \frac{\bar{c}_{1j}^2 \text{Max}(\bar{A}_j)}{ng(\bar{c}_{ij}^2)} - \text{Max}(\bar{b}_j), \right. \\ \left. \sum_{j=1}^n \frac{\bar{c}_{1j}^2 \text{Std}(\bar{A}_j)}{ng(\bar{c}_{ij}^2)} - \text{Std}(\bar{b}_j), \sum_{j=1}^n \frac{\bar{c}_{1j}^2 \text{Median}(\bar{A}_j)}{ng(\bar{c}_{ij}^2)} - \text{Median}(\bar{b}_j) \right\},$$

where $\bar{c}_{ij}^2 = \bar{c}_{ij} \bar{c}_{ij}$ for each $i \in \{1, \dots, p\}$ and $j \in \{1, \dots, n\}$.