

Highlights

Optimization with Constraint Learning: A Framework and Survey

Adejuyigbe Fajemisin, Donato Maragno, Dick den Hertog

- We provide a framework for optimization with constraint learning.
- We review the literature using constraint learning in light of our proposed framework.
- Current trends are discussed and opportunities for future work are highlighted.

Optimization with Constraint Learning: A Framework and Survey

Adejuyigbe Fajemisin, Donato Maragno, Dick den Hertog

Amsterdam Business School, University of Amsterdam, 1018 TV Amsterdam, The Netherlands

Abstract

Many real-life optimization problems frequently contain one or more constraints or objectives for which there are no explicit formulas. If data is however available, these data can be used to learn the constraints. The benefits of this approach are clearly seen, however there is a need for this process to be carried out in a structured manner. This paper therefore provides a framework for Optimization with Constraint Learning (OCL) which we believe will help to formalize and direct the process of learning constraints from data. This framework includes the following steps: (i) setup of the conceptual optimization model, (ii) data gathering and preprocessing, (iii) selection and training of predictive models, (iv) resolution of the optimization model, and (v) verification and improvement of the optimization model. We then review the recent OCL literature in light of this framework, and highlight current trends, as well as areas for future research.

Keywords: Analytics, Optimization, Constraint learning, Machine learning

Email addresses: a.o.fajemisin@uva.nl (Adejuyigbe Fajemisin),
d.maragno@uva.nl (Donato Maragno), d.denhertog@uva.nl (Dick den Hertog)

Preprint submitted to European Journal of Operational Research *September 30, 2021*

1. Introduction

1.1. *The Synergy between Optimization and Machine Learning*

1.1.1. *Optimization for Machine Learning*

Many problems in Machine Learning (ML) can be written as optimization problems (Bennett and Parrado-Hernández, 2006). In ML, a loss function provides a measure of the difference between the value predicted by the ML algorithm and the actual value. It follows then that ML practitioners seek to minimize this loss function, which leads to the use of Optimization for ML. In certain applications, e.g. deep neural networks, the objective function of the learning task is non-convex, leading to the application of non-convex optimization methods (Jain and Kar, 2017). The problem of selecting the right parameters for an ML model to give the best performance is called hyperparameter tuning. This can be accomplished using heuristic methods such as random search or grid search (Bergstra et al., 2011). Nevertheless, this problem also lends itself to more traditional optimization methods such as gradient descent (Yang and Shami, 2020). In the book (Bertsimas and Dunn, 2019), one of the key underlying ideas is that machine learning is inherently linked to optimization, rather than probability theory. They demonstrate this by using mixed-integer optimization (MIO) techniques to construct optimal decision trees, and show that their MIO-constructed trees outperform classification and regression trees (CART) (Breiman et al., 1984) (which are constructed using heuristic methods) by up to 7% in some cases. Clearly there is a huge scope for optimization in ML. The paper (Sun et al., 2019) provides a recent survey of optimization methods from an ML perspective.

1.1.2. *Machine Learning for Optimization*

In the optimization field, there is more and more attention for using ML for optimization. Optimization problems are prevalent in all sorts of industries, including logistics, agriculture, energy, and many others. In practice, these problems can be large and difficult to solve. It stands to reason therefore that given the recent advances in ML (e.g. (Silver et al., 2016)), ML is increasingly finding a role in combinatorial optimization (CO). There are more and more works exploiting the flexibility and usefulness of ML to aid in the solution of difficult optimization problems. In (Abbasi et al., 2020) for example, the authors train ML models on the optimal solutions of a large-scale supply chain problem to learn optimal decisions. This allows the ML model can be used on a daily basis, instead of needing to solve the large-scale

optimization problem. In another example, Fischetti and Fraccaro (2019) use ML and optimization to speed up the solution of a wind farm layout optimization problem. They use solutions from already solved problems to train ML models, with the aim of being able to predict the solutions of a new layout problem without lengthy computation times. In addition to speeding up the solution of problems, ML and deep learning have also been used to speed up tree search in (Hottung et al., 2020) for example, to accelerate the Branch-and-Price Algorithm in (Vaclavik et al., 2018), and to guide branching in Mixed Integer Programs (MIP) (Khalil et al., 2016). A recently published survey (Bengio et al., 2021) provides an overview of recent attempts to leverage machine learning to solve combinatorial optimization problems.

1.2. Motivation for this Study

The examples in the previous paragraph use ML to learn or obtain solutions to optimization problems in a more efficient manner. In our view however, there is even more value in using ML for optimization: ML can be used to learn constraints. Many real-life optimization problems contain one or more constraints or objectives for which there are no explicit formulae. However when there is data available, one can use the data to learn the constraints. Two such examples are given below.

Example 1 - The World Food Programme (WFP): This is the problem of providing palatable rations to recipients of the United Nations (UN) WFP. The paper (Peters et al., 2021), contains a MIP model for optimizing the food supply chain for the WFP. A key part of this model is the diet model (Stigler, 1945; Dantzig, 1990), which is concerned with selecting a set of foods to satisfy a set of daily nutritional requirements at minimum cost. Besides providing food, it is of additional benefit to the recipients that the provided rations be palatable. These palatability constraints cannot be manually developed. However, abundant data has been used to learn the palatability constraints (Maragno et al., 2022), which can then be added to the model.

Example 2 - Radiotherapy Optimization: Optimization is used nowadays to optimize treatment plans for thousands of cancer patients every day (Eikelder et al., 2019). However, e.g., the nonlinear formula for the Tumour Control Probability (TCP) is based on a population average. Together with

researchers from Harvard Medical School, we are in the process of using the abundant patient data that is available, and learning the TCP constraint function including patient characteristics. This learned constraint can then be added to the final optimization model.

In addition to using traditional ML approaches to learn constraints for optimization problems (e.g. regression trees and linear regression in (Verwer et al., 2017) and neural networks in (Villarrubia et al., 2018)), other techniques such as genetic programming (e.g. in (Pawlak and Krawiec, 2018)), mixed-integer linear programming (MILP) (e.g. in (Pawlak and Krawiec, 2017a)) and symbolic regression have also been used. In our literature search, we have noticed several works of this kind from several different application areas and even several scientific fields. However it seems beneficial to have these works collated and discussed in one place. Additionally, we propose to take a more holistic view of the entire constraint learning process, starting from data gathering and processing, to resolution and validation of the final optimization model.

This paper is therefore intended to provide both a survey of literature on constraint learning for optimization, as well as offer a framework for effectively facilitating Optimization with Constraint Learning (OCL). The five steps of the proposed framework are: (i) setup of the conceptual optimization model, (ii) data gathering and processing, (iii) selection and training of the predictive model(s), (iv) resolution of the optimization model, and (v) validation/improvement of the optimization model. These steps will be discussed in more detail throughout this paper. We use the term “Constraint Learning” to make clear that we are interested in cases in which most of the optimization model is known, but some of the constraints are difficult to model explicitly. This is to differentiate from other terms such as “model learning” (Donti et al., 2017) and “model seeking” (Beldiceanu and Simonis, 2016) seen in the literature. We should also note that there are several approaches in the literature which are similar to *constraint learning* but do not fit in this survey for several reasons. In Constraint Programming (CP) for example, the process of learning CP models from data (positive and/or negative examples) has been widely studied in papers such as (Bessiere et al., 2004; O’Sullivan, 2010; Beldiceanu and Simonis, 2016; Kolb, 2016; Galassi et al., 2018). These are not included as we are interested in cases where predictive models are embedded as constraints in mathematical optimization. Additionally, papers such as (Bertsimas and Kallus, 2020; Elmachtoub and

Grigas, 0; Demirovic et al., 2019) which focus on predicting problem *parameters*, or use predictive models to approximate *known* objective functions and constraints (such as (Villarrubia et al., 2018)) are excluded. We also exclude papers such as (Nascimento et al., 2000; Nagata and Chu, 2003; Venzke et al., 2020a; Venzke and Chatzivasileiadis, 2020) where the structure of the problem is known, but predictive models are used to replace the problem in order to speed up computation. Similarly, papers like Shang et al. (2017); Han et al. (2021) which use predictive models to learn the uncertainty set for robust optimization problems are not considered. In Gilan and Dilkina (2015); Bagloee et al. (2018); Say and Sanner (2018); Say et al. (2020), although there is an interplay between ML and optimization in some way, these papers are out of the scope of this survey as they do not focus on constraint learning.

While Lombardi and Milano (2018) and Kotary et al. (2021) provide short surveys on this subject, and we present a more extensive framework for constraint learning and consider a wider range of the literature. Each paper identified in our literature search will be discussed in light of the proposed framework, and gaps in knowledge and scope for future lines of research will be identified. Running examples will be used to illustrate the steps of the framework. It is envisaged that in addition to providing a review as well as the framework, our work will bring together ideas from diverse fields.

1.3. Outline

This framework and survey is structured as follows. In Section 2, we first describe the running examples that will be used to illustrate our framework. Then in Section 3, we describe this framework and review the recent literature on light of it in Section 4. Areas for further research are presented in Section 5 with general conclusions given in Section 6.

2. Running Examples

In order to illustrate the concepts presented in later sections, we first present two practical optimization problems in which constraint learning plays an important role. These examples have been mentioned in Section 1.2 but are now given in more detail.

2.1. Palatability Considerations in WFP Food Baskets

This first example deals with the issue of providing palatable rations in the food baskets delivered by the WFP to population groups in need.

These food baskets are designed to meet the nutritional requirements of a population, and take into account existing levels of malnutrition and disease, climatic conditions, activity levels, and so on (UNWFP, 2021). In addition to considering the nutritional content, we would also like to be able to consider the palatability of the rations. This problem is particularly relevant in that the palatability of a food basket can change drastically between different groups of beneficiaries. Since the palatability constraint is difficult to model manually, and a data set with different food baskets and their palatability scores is available, a predictive model can be used to derive the palatability constraint from the data, which can then be embedded into the original food basket optimization problem.

Let \mathcal{K} be the set of commodities (such as rice, flour, chickpeas, sugar, etc.) which may be included in a food basket, and \mathcal{L} be the set of nutrients (such as vitamins, minerals, proteins, etc.) which are important for maintaining good health. If the nutritional requirement for nutrient $l \in \mathcal{L}$ (in grams/person/day) is denoted by $Nutreq_l$, and the nutritional value for nutrient l (per grams of commodity $k \in \mathcal{K}$) is denoted by $Nutval_{kl}$, then a simple model for food basket optimization problem can be given as

$$\min_{\mathbf{x}} \mathbf{c}^T \mathbf{x} \tag{1a}$$

$$\text{s.t.} \quad \sum_{k \in \mathcal{K}} Nutval_{kl} x_k \geq Nutreq_l \quad \forall l \in \mathcal{L} \tag{1b}$$

$$Palatability_score(\mathbf{x}) \geq P_{min} \tag{1c}$$

$$x_k \geq 0 \quad \forall k \in \mathcal{K}, \tag{1d}$$

where $\mathbf{x} \in \mathbb{R}^{|\mathcal{K}|}$ is the vector of commodities, x_k is the amount of each commodity $k \in \mathcal{K}$ (measured in grams) included in the food ration, $\mathbf{c} \in \mathbb{R}^{|\mathcal{K}|}$ is the vector of costs associated with each commodity (in \$/gram), and the objective function (1a) is a minimization of the ration cost. Constraint (1b) is used to ensure a minimum level of nutritional requirements in the rations. Constraint (1c) is the embedding of the predictive model used to ensure that the palatability of the ration is at least equal to P_{min} . The diet problem in (1a) to (1d) is only a part of a much bigger model that also includes the sourcing, facility location and transportation problems. The constraint to be learned here is a function of relatively few variables ($|\mathcal{K}|$ in this case), compared with the thousands of variables in the full model.

The predictive model is trained on a sparse data set \mathbf{D} in $\mathbb{R}^{n \times (|\mathcal{K}|+1)}$ composed of n samples, with each sample composed of $|\mathcal{K}|$ features and one label. Each sample is a different food basket where the $|\mathcal{K}|$ features are the amount of each commodity k that make up the food basket, and the label corresponds to the palatability score associated with each specific food basket. This label can either be numeric or categorical, resulting in $Palatability_score(\mathbf{x})$ being either a regression or classification model respectively.

2.2. Radiotherapy Optimization

The second example problem we will use to illustrate the framework is radiotherapy (RT) optimization. In RT, ionizing radiation beams are used to control tumor growth by damaging cancerous tissue. Healthy, non-cancerous tissue (also known as organs-at-risk (OAR)) are however present in close proximity to the tumors, and consequently incur unavoidable damage. The tumour damage is quantified as the Tumor Control Probability (TCP), while the unavoidable damage to the OARs is quantified using the Normal Tissue Complications Probability (NTCP). The final optimization problem is to maximize the TCP subject to a set of constraints ensuring that the NTCP does not exceed a specified level.

The radiation beams are delivered from a radiation source such as a linear accelerator or a cyclotron, and directed towards the tumor cells from multiple angles using a rotating gantry. The angles of delivery are predetermined by clinicians and at each given angle, the radiation beam is shaped by passing it through a multileaf collimator (MLC) (Figure 1a). The MLC adjusts the shape of the radiation beam so that it generally resembles the shape of the tumor (Figure 1b).

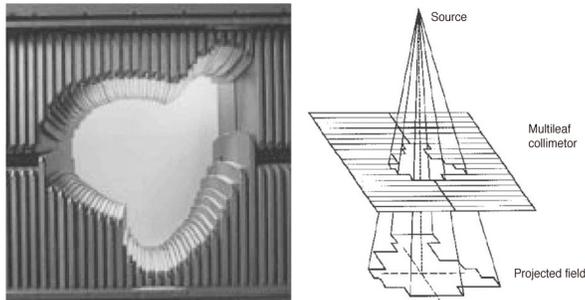


Figure 1: Radiation treatment delivery: (a) MLC (b) MLC impact on beam shape (Kim, 2008)

To facilitate the treatment planning process, the target volumes (both tumors and OARs) are discretized into 3-dimensional volumetric pixels (or voxels), while each radiation beam is discretized into smaller sub-units called beamlets. Let n be the total number of available voxels, and \mathbf{x} be the m -dimensional vector of beamlet intensities (also known as a fluence map). Then, the dose deposition matrix \mathbf{A} is an $n \times m$ -dimensional matrix, where each value A_{ij} determines the unit intensity of dose delivered to voxel i from beamlet j . The physical radiation dose \mathbf{d} delivered to each voxel i is an n -dimensional vector, and is the product of \mathbf{A} and \mathbf{x} . A schematic diagram of the radiation delivery is shown in Figure 2.

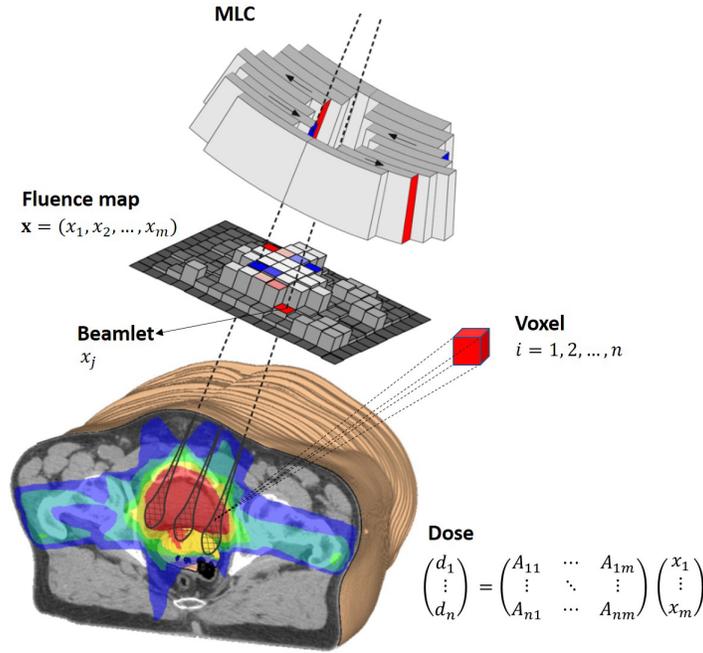


Figure 2: Schematic diagram of radiation delivery (Müller, 2009)

If \mathcal{Q} is the set of OARs (e.g. stomach, liver, spinal cord) which need to

be protected, then the optimization problem can be written as

$$\max_{\mathbf{x}, \mathbf{d}} \text{TCP}(\mathbf{d}, \boldsymbol{\gamma}) \quad (2a)$$

$$\text{s.t. } \text{NTCP}_q(\mathbf{d}, \boldsymbol{\gamma}) \leq \delta_q \quad \forall q \in \mathcal{Q} \quad (2b)$$

$$f(\mathbf{d}) \leq 0 \quad (2c)$$

$$\mathbf{Ax} = \mathbf{d} \quad (2d)$$

$$\mathbf{x} \geq 0, \quad (2e)$$

where $\boldsymbol{\gamma}$ is a vector containing the biological characteristics of the tumor and the OARs, as well as other patient-specific information. δ_q is the maximum allowable damage to OAR q while equation (2c) is a placeholder for other constraints on the dose distribution \mathbf{d} (e.g. minimum dose, maximum dose, etc.) and is usually convex or can be converted to a convex form. The decision variables \mathbf{x} and \mathbf{d} are the set of the radiation intensities and the doses to be delivered over the course of the treatment respectively. The problem size depends on the number of beamlets, constraints, as well as distinct OARs considered in the treatment planning. A more detailed treatment of intensity-modulated radiation therapy optimization can be found in (Hoffmann et al., 2008).

In the current literature, explicit expressions are used for the TCP and NTCP functions (Gay and Niemierko, 2007). These however are surrogate functions based on population-wide estimates, and as such the response of different patients to the same treatment can be quite varied (Scott et al., 2021). Thus, while some patients experience complete remission with minimal damage to OARs, others suffer from radiation poisoning or tumor recurrence. It has been shown in (Tucker et al., 2013) that the inclusion of patient-specific characteristics (such as genetic predispositions) can improve the prediction of patients' response to RT treatment. We therefore intend to use data to learn (2a) and (2b), leading to the creation of more effective treatment plans for patients.

The above two examples have illustrated the potential benefits of OCL. In the next section, we will describe the general framework for OCL, as well as show how to apply the steps of the proposed framework to optimization problems.

3. General Description of Framework

The proposed framework is the result of a careful analysis of many approaches adopted in the literature covered by this survey. It is the formalization of a procedure that (partially) recurs in most of the approaches for learning constraints. We believe that a complete view of all the steps that characterize the framework will help future research see the big picture of constraint learning, and the impact that each step has on the final performance (in a loose sense). These steps are enumerated below:

1. Setup of the conceptual optimization model
2. Data gathering and preprocessing
3. Selection and training of predictive models
4. Resolution of the optimization model
5. Verification and improvement of the optimization model.

Although the steps of the framework are given in sequential order, it should be noted that some of these steps may be repeated in non-sequential order, i.e. internal loops are possible. This may be necessary in order to improve the performance of the optimized model.

3.1. Step 1: Setup of the Conceptual Optimization Model

A graphical overview of this first step is given in Figure 3. It consists of five sub-steps, of which the first three are also commonly performed in traditional operations research (OR) modeling. This initial stage is usually done in collaboration with domain experts. This step begins with (i) defining the decision variables and parameters involved, and (ii) defining those constraints that are easy to model manually.

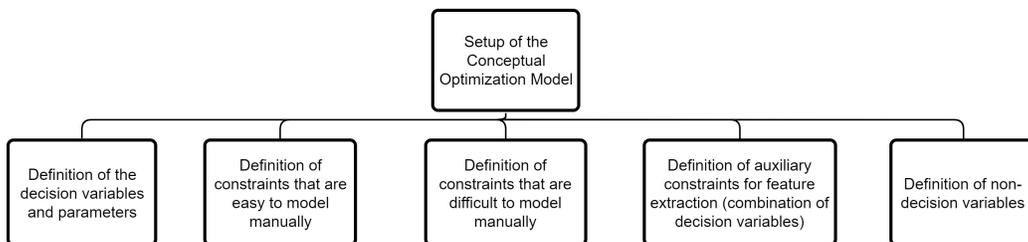


Figure 3: Graphical overview of Step 1: Setup of the conceptual optimization model

Mathematically, let $\mathbf{x} = (x_1, \dots, x_n)$ be the vector of decision variables for an optimization problem which may be written as

$$\min_{\mathbf{x}} f(\mathbf{x}) \tag{3a}$$

$$\text{s.t. } \mathbf{g}(\mathbf{x}) \leq 0 \tag{3b}$$

$$\mathbf{x} \in X, \tag{3c}$$

where the objective function $f(\mathbf{x})$ is being minimized (and may also be maximized without loss of generality), $\mathbf{g}(\mathbf{x})$ are known constraints of the problem, and X is a bound constraint on \mathbf{x} . The constraints $\mathbf{g}(\mathbf{x})$ are well-defined and represent the case in which a process or behaviour can be modelled manually.

In certain cases however, some constraints are difficult to model manually or to state in concise mathematical terms. A traditional approach would suggest approximating these complex constraints with other simpler constraints, but this results in a loss of model accuracy. By using constraint learning methods however, predictive models can be used to design data-centric constraints.

Let \mathbf{y} represent a quantity of interest which is difficult to model. If \mathbf{y} is a function of the decision variables \mathbf{x} , then a predictive model can be used to learn \mathbf{y} :

$$\mathbf{y} = \hat{\mathbf{h}}(\mathbf{x}). \tag{4}$$

This learned function can then be embedded as a constraint into the original problem. The palatability constraint of Section 2.1 is an example of a quantity of interest which is difficult to model. \mathbf{y} could be the palatability, and is a function of the amount of each commodity x_k in the food basket. It should be noted that, depending on the predictive model learned, it might not be possible for optimization solvers to directly handle the learned models. $\hat{\mathbf{h}}(\mathbf{x})$ is therefore more likely to be an encoding of the predictive model, and not the actual model. If $\boldsymbol{\theta}(\mathbf{y}) \leq 0$ are constraints on \mathbf{y} , then (3a) - (3c) become

$$\min_{\mathbf{x}, \mathbf{y}} f(\mathbf{x}) \tag{5a}$$

$$\text{s.t. } \mathbf{g}(\mathbf{x}) \leq 0 \tag{5b}$$

$$\mathbf{y} = \hat{\mathbf{h}}(\mathbf{x}) \tag{5c}$$

$$\boldsymbol{\theta}(\mathbf{y}) \leq 0 \tag{5d}$$

$$\mathbf{x} \in X. \tag{5e}$$

It should be noted that some part of the known constraint might also be a function of \mathbf{y} , i.e. $\mathbf{g}(\mathbf{x})$ could be $\mathbf{g}(\mathbf{x}, \mathbf{y})$, however we will continue with the definition where the known and unknown constraints are written separately. Similarly, some part of the objective function $f(\mathbf{x})$ might also depend on the quantity to be learned (i.e. $f(\mathbf{x}, \mathbf{y})$). The objective function $f(\mathbf{x})$ can also be learned, but we will also continue with the definition where $f(\mathbf{x})$ is known. We would also like to make a subtle distinction between *constraint learning* and *constraint function learning*. In the former case, we learn a constraint using its feasibility threshold during the model training process. In the latter case, we learn a continuous predictive function without considering the feasibility threshold, and we add this threshold as a constraint in the optimization model. In either case, the above formulation holds.

One of the factors that affects the performance of a predictive model is the quality of the features used to find patterns in the dataset. In (5c), the features used in the predictive model correspond to the vector of decision variables \mathbf{x} (or a subset of \mathbf{x}). This is not always the best set of features. In fact, the predictive model's performance can be often improved with the addition of other features \mathbf{w} that are not part of the decision-making process, but convey additional information about the problem. For example in Section 2.2, although the decision variables are radiation intensities \mathbf{x} , additional information is conveyed by the vector $\boldsymbol{\gamma}$ which contains patient-specific information. The inclusion of these variables helps to improve the predictive model's performance, and therefore the performance of the optimization model. Replacing (5c) with $\mathbf{y} = \hat{\mathbf{h}}(\mathbf{x}, \mathbf{w})$ gives

$$\min_{\mathbf{x}, \mathbf{y}} f(\mathbf{x}) \tag{6a}$$

$$\text{s.t. } \mathbf{g}(\mathbf{x}) \leq 0 \tag{6b}$$

$$\mathbf{y} = \hat{\mathbf{h}}(\mathbf{x}, \mathbf{w}) \tag{6c}$$

$$\boldsymbol{\theta}(\mathbf{y}) \leq 0 \tag{6d}$$

$$\mathbf{x} \in X. \tag{6e}$$

Although the \mathbf{w} 's are known in the radiotherapy optimization example (e.g. patient-specific information such as gene expressions), in certain cases they may not be known and as such, we may have to optimize for the average or worst case. It may also be possible to treat any uncertainties in the \mathbf{w} 's using a robust or stochastic optimization approach.

The vector \mathbf{x} can also be exploited to build more complex features using a process known as *feature engineering*. The idea behind this process is to use effective combinations of decision variables as new features to improve the performance of the predictive model, reduce the number of features used in the learning process, or embed prior knowledge into the predictive model. These new features $\mathbf{z} = (z_1, \dots, z_m)$ can be expressed using auxiliary constraints $s_i(\mathbf{x})$, where s_i is the specific function used to extract the feature z_i :

$$\min_{\mathbf{x}, \mathbf{y}} f(\mathbf{x}) \tag{7a}$$

$$\text{s.t. } \mathbf{g}(\mathbf{x}) \leq 0 \tag{7b}$$

$$\mathbf{y} = \hat{\mathbf{h}}(\mathbf{x}, \mathbf{w}, \mathbf{z}) \tag{7c}$$

$$\boldsymbol{\theta}(\mathbf{y}) \leq 0 \tag{7d}$$

$$z_i = s_i(\mathbf{x}) \quad i = 1, \dots, m \tag{7e}$$

$$\mathbf{x} \in X. \tag{7f}$$

The palatability constraint in Section 2.1 is an example of such a case. It is not simply a function of the decision variables, but more accurately a function of aggregates of \mathbf{x} . These z_i variables represent the amount of food per macro category (Cereals & Grains, Pulses & Vegetables, Oils & Fats, Mixed & Blended foods, Meat & Fish & Dairy), so that (1c) becomes *Palatability_score*(\mathbf{z}) $\geq P_{min}$.

Model (7a) - (7f) is a broad formulation of all the fundamental aspects that characterize optimization with constraint learning. \mathbf{x} is the vector of decision variables, $\mathbf{g}(\mathbf{x})$ are constraints which are easy to model, and \mathbf{y} is the difficult part of the model to design manually. It is the output of a predictive model $\hat{\mathbf{h}}(\cdot)$ which takes as input decision variables \mathbf{x} , non-decision variables \mathbf{w} , and \mathbf{z} , which is the vector of variables designed as a combination of the \mathbf{x} 's.

3.1.1. Trust Regions

The performance of a predictive model worsens when it is evaluated on data that is much different from the data used to train it. During the optimization process, solutions far from what has been seen in the training and test sets could be evaluated. In this case, the predictive performance could

differ much from our expectations. In order to address this issue, we can restrict the optimization to a region that is densely populated by samples that have been used to train and test the predictive model. This is known as a Trust Region (TR), and we assume that within this region, the predictive performance will be consistent with what we have seen during the model evaluation. The concept of a trust region is used in (Biggs et al., 2018) and is thoroughly explained in (Maragno et al., 2022) where the authors constrain the optimal solution to be within the convex hull of all samples (training and test samples regardless of their label). This prevents the optimization process from evaluating solutions very different from what has been seen during the training and evaluation of the model. A larger TR might lead to a better solution in terms of objective function value, but also to less confidence in the performance score obtained in the predictive model evaluation.

3.2. Step 2: Data Gathering and Preprocessing

Once the conceptual optimization model has been formulated, the next step is to focus on collecting and preprocessing the data which will be used to learn the constraints (7c). As data gathering and processing is a huge subject in its own right, we will only discuss those aspects that are important in the context of constraint learning. Unlike the first step of the proposed framework, there is no order to the different sub-steps here; we simply highlight the different characteristics of the data, and how best to exploit them for our purposes. A graphical overview of this second step is given in Figure 4.

These characteristics inform the tasks to be performed, and have a direct impact not only on the predictive model, but also on the final optimization model. For example, the amount of data available might influence the choice of the ML method used to learn the constraint, which may then have an influence on the approach used to solve the final optimization model.

Understanding the data required to train and test predictive models is an activity that requires the support of domain experts. The data may already be available or may need to be generated or collected. In the second case, the best way to generate or collect and store data must be determined, and the data may either be obtained from real-world processes, or obtained from simulations of the processes. Whether or not the data is real or simulated depends on the problem under consideration, and the availability of resources (in terms of time and money). If the right data are available, then they can be used directly (with or without preprocessing) to learn the constraints (7c). If the data are not yet available, decisions have to be made on how to collect

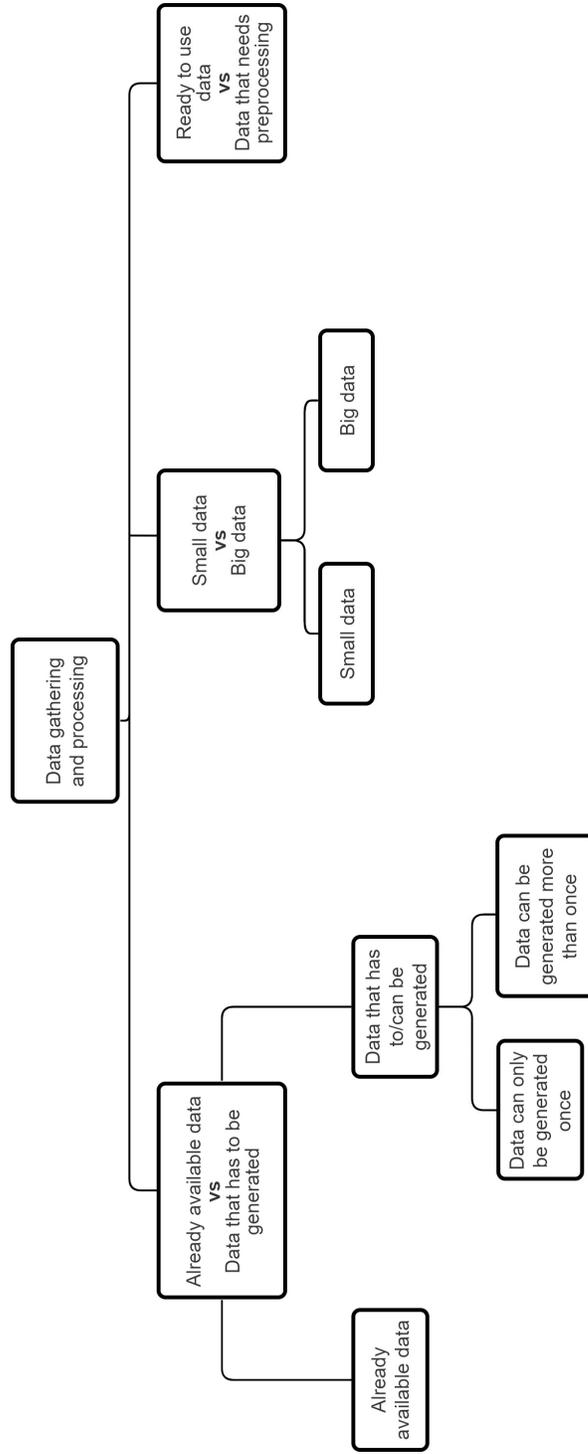


Figure 4: Graphical overview of Step 2: Data gathering and preprocessing

the data. Data collection or generation strategies should follow established principles such as Design of Experiments (DoE) techniques (which are applicable for physical experiments) or Design of Computer Experiments (DoCE) techniques (applicable for computational experiments). In the WFP application of Section 2.1 for example, collecting data on the palatability of rations during an epidemic or in a war zone might not be possible. Generation or simulation of data following DoCE principles is useful in such a situation. A recent review of DoCE is given in (Garud et al., 2017).

Data collection or generation can take place either in a single run or in multiple runs. Performing data collection multiple times can make the learning and optimization process more accurate. For example if the simulation of a process from which data is obtained takes several hours, then the choice of input parameters to the simulator might need to be carefully considered. Adaptive DoCE techniques such as those in (Crombecq et al., 2011) and (Xu et al., 2014) may be used in such cases. In Derivative Free Optimization (DFO) only a few or even one new data point is collected in each iteration to minimize the number of data points needed to obtain an optimal solution. We refer to Conn et al. (2009) for a detailed treatment on DFO.

In addition to the availability or sourcing of the data, the size of the dataset also informs the choice of the learning approach. It is known that deep neural networks are able to exploit large datasets better than traditional algorithms such as logistic regression or support vector machine (Ng, 2020). In contrast, methods such as kriging (Matheron, 1963) can be used with smaller datasets.

Regardless of the size of the datasets, most data (especially those not generated with the use of simulators) need to be pre-processed before they can be used in the next step (learning). Many preprocessing activities have been widely discussed in ML literature but common preprocessing steps include imputation, normalization or scaling, binarization and binning (Pedregosa et al., 2011). Imputation is the process where missing values in a dataset are replaced by estimates, while normalisation or scaling is where the values of the features are scaled to an interval (usually $[-1, 1]$ or $[0, 1]$) to avoid numerical ill-conditioning (Santner et al., 2003). Binarization is the process of converting data to binary values based on whether they are above a threshold or not, and binning involves creating new features as discretized versions of existing continuous variables. Whatever data collection, generation and preprocessing approaches are used, the end result of this second step is a dataset that is ready to be used for training, testing and validating

the predictive algorithms. A review of the various data generation and pre-processing approaches used in the context of constraint learning is given in Section 4.2.

3.3. Step 3: Selection and Training of Predictive Model(s)

Once the data gathering and preprocessing activities have been completed, it is necessary to select, train, validate, and test the predictive model that will subsequently be used to represent one or more constraints in the final optimization model. Selection and training are two activities that go hand in hand. It is not possible to select the best predictive model without first training and testing it on the available dataset. Correspondingly, poor performance of a model after training might require a different predictive model to be selected. A graphical overview of this third step is shown in Figure 5.

The selection of the predictive model is based on six main factors:

- **Classification or regression:** The choice of predictive model will depend on whether the constraint to be learned results in a discrete set of values (classification), or a continuous spectrum of values (regression). This has consequences on the format of the predictive model output, thus limiting the choice to a restricted set of models. In constraint learning, the use of a classification model results in a feasible set being learned, while the use of a regression model results in a function being learned. The palatability score function in (1c) can for example, be a regression model, where the output is a continuous value giving a measure of the diet's palatability. It can also be a classification function where the output is a binary variable stating whether the diet is palatable or not.
- **Computational complexity:** The choice of predictive model directly affects the computational complexity of the final optimization model. In contrast to a linear regression model, deep learning models require several mathematical or logical statements to represent them and undoubtedly increase the complexity of the final optimization problem. This increased complexity may compromise the optimality of the final solution, and a trade-off may be required. For example, fewer layers may be used, or the number of neurons in a layer may be small. A decision might also be made to select a predictive model so that the complexity of the final optimization problem does not increase. In the

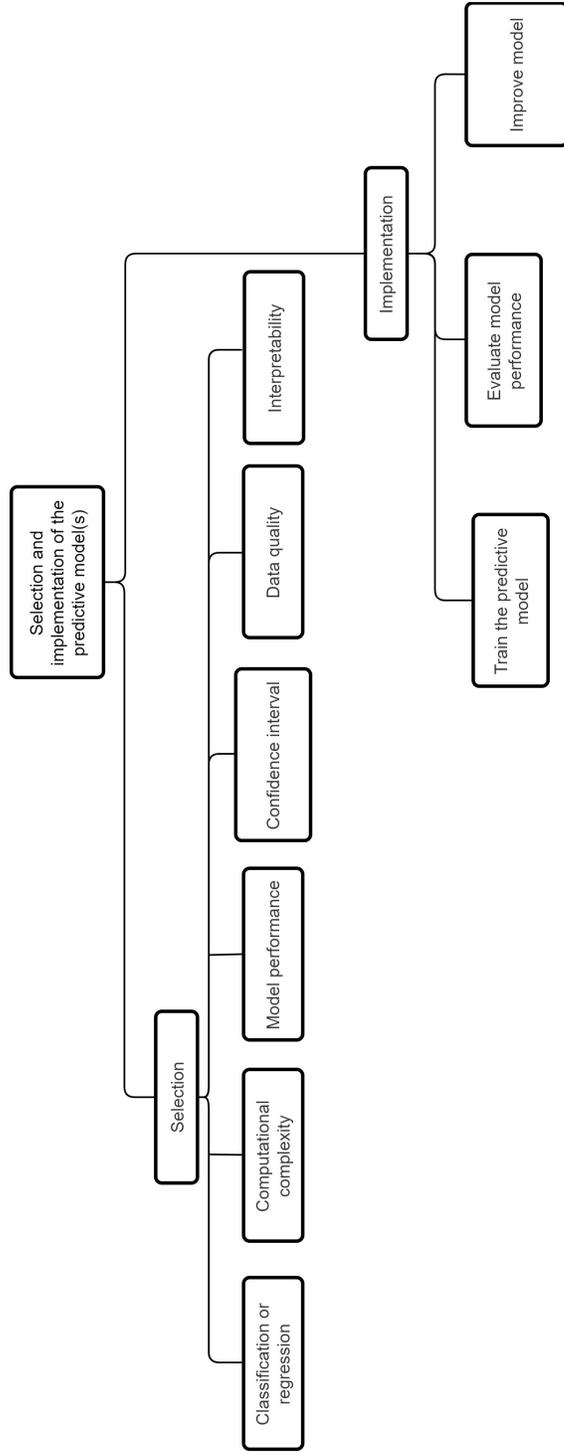


Figure 5: Graphical overview of Step 3: Selection and training of predictive models

WFP example, the underlying problem is linear so a linear predictive model may be chosen to learn $Palatability_score(\mathbf{x})$.

- Model performance: In addition to computational complexity, the predictive model has to perform adequately with respect to some performance measure (e.g. accuracy or mean squared error). A simple but inaccurate model can have worse consequences than a complex but accurate model. The performance should be evaluated on both a test set and in the final optimization model. Predictive models that perform well on the test set may perform poorly on the “extreme” solutions evaluated by the solver during the optimization process, and so this means that the performance evaluation has to be performed both in Step 3 and Step 5.
- Confidence intervals: If the predictive models used provide confidence intervals, it might be possible to use this information in approaches such as stochastic optimization or robust optimization.
- Data quality: Despite the attention focused on the dataset in Step 2, it may be necessary to choose a predictive algorithm that is robust to noise or missing data. This is especially the case if the data to be used has been provided by a third-party.
- Interpretability: The selection of a predictive model might also depend on whether or not the practitioner requires the model to be interpretable. In the radiotherapy optimization problem of Section 2.2, clinicians may prefer highly explainable predictive models to be used to learn the $TCP(\mathbf{d}, \boldsymbol{\gamma})$ and $NTCP_q(\mathbf{d}, \boldsymbol{\gamma})$ functions due to the seriousness of the application. There are however pros and cons to consider. Decision trees, for example, are known to be highly explainable but may require the linearization of several disjunctions to encode them in the optimization problem.

A review of the various predictive models used for constraint learning in optimization is given in Section 4.3.

The implementation process is a cycle of training the predictive model, evaluating the model’s performance, and improving the model in terms of the data, model parameters, and even the model used. If a feed-forward neural

network is used for example, parameters such as the number of layers, the number of nodes per layer, the activation functions, whether or not to use dropout regularization, etc., need to be determined. Out-of-sample testing should be done to ensure adequate performance of the model, with changes made to improve performance if necessary.

In the context of constraint learning, under the same computational complexity, the best predictive model is the one with higher performance in the final optimization model. The goal should be to generate prediction models that aim to minimize decision error, and not just prediction error.

3.4. Step 4: Resolution of the Optimization Model

Once the predictive model has been used to learn the difficult-to-model constraint, it will need to be embedded into the rest of the optimization problem. Figure 6 gives a graphical overview of this step. Recall that the predictive model $\mathbf{h}(\cdot)$ in (7c) may not be easy to add as a constraint, and may need to be encoded in a such a way that it can be handled by commercial or conventional solvers. This however depends on the type of model used, as a linear or quadratic function can be embedded more easily than a neural network for example. The integration of the learned constraints is often not trivial, and may result in a significant reformulation of the model, while trying not to affect its predictive performance. If possible, the predictive model may be incorporated either by providing gradients directly to the solver, or by reformulating it such that it becomes linear, convex, conic, etc., and can more readily be incorporated into the optimization problem. If desired, the optimization models may be modified to take into account any uncertainty measures provided by the predictive models. A deeper look at the various methods used in the literature to encode predictive models is given in Section 4.4.

In addition to encoding the learned constraints in a manner that the solver can exploit, a successful embedding can also include providing additional information (such as bounds) to the solver of choice in order to facilitate a more efficient solution process. Not only is it necessary to encode the learned constraints in a manner that the solver can handle, a successful embedding should also make it possible to take full advantage of the optimization solver’s capabilities.

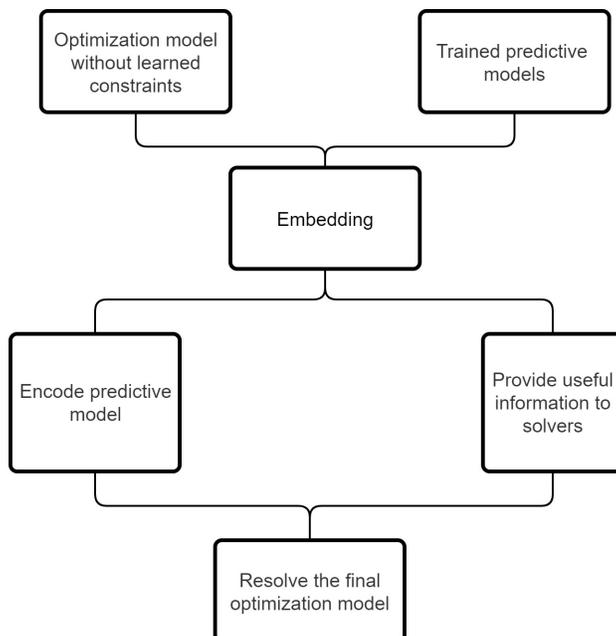


Figure 6: Graphical overview of Step 4: Resolution of the optimization model

3.5. Step 5: Verification and Improvement of the Optimization Model

In this final step illustrated in Figure 7, we evaluate the performance of the constraint learning process with respect to three things: the final optimization model, the predictive model, and the data. For the optimization model, the focus is on both the goodness of the solution (i.e. the objective function value), as well as the time taken to find this solution. For the performance of the predictive model, the focus is on ensuring that the learned constraints accurately represent the difficult-to-model constraints. If it is possible to generate data, then using different data to improve the predictive model can in turn help to improve the optimal solution. One way of doing this is via a process called active learning. Here, the dataset is continually modified by adding new solutions (output of the optimization model) to it, which is then used to re-train the predictive model. A review of the approaches to used to improve verify and improve models with learned constraints is provided in Section 4.5.

While the above framework steps have been given in sequential order, it should be reiterated that some of these steps may be repeated. For example,

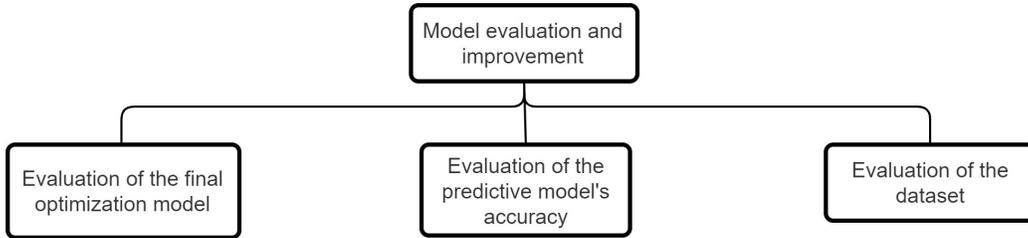


Figure 7: Graphical overview of Step 5: Verification and improvement of the optimization model

if it is seen that in step 5 that the learned constraint is a poor representation of the reality, we may need to go back to step 1. Similarly, the incorporation of new data will mean that the predictive model needs to be retrained and re-embedded, and so on. Internal loops are also possible — it may be necessary to iterate between steps 2 and 3 until an acceptable predictive model is developed, or between steps 4 and 5 until satisfactory results are achieved.

4. Review

4.1. Setup of the Conceptual Model

In the literature, only a few papers have attempted to formalize the process of optimization with constraint learning. The papers by Lombardi et al. (2017) and Maragno et al. (2022) provide the most detailed approaches at formalising this process. There is also a clear structure provided in (De Angelis et al., 2003) for learning unknown constraints. The methodology in Lombardi et al. (2017), called Empirical Model Learning (EML), is described as a way of obtaining components of a prescriptive model from machine learning. Similar to the steps given in Figure 3, they define a conceptual model including known constraints, constraints to be learned, as well as the definition of constraints used for feature extraction. Maragno et al. (2022) also provide a general framework for OCL that can be applied every time data is available. A key focus of this approach is to learn constraints and embed them in such a way that the computational complexity of the final model is not increased. Padmanabhan and Tuzhilin (2003) and Deng et al. (2018) also conceptually discuss the use of data mining to define objectives and constraints.

While other papers which use a predictive element to learn constraints do not have as detailed a setup, they incorporate some of the elements of Step 1

in their approach. Some authors (e.g. (Verwer et al., 2017)) explicitly define constraints to perform feature extraction, while most of the papers reviewed simply identify which constraints are difficult to model manually and replace them with predictive models. This will be seen later in Sections 4.3 and 4.4.

The approach taken to formalize the constraint learning process in Pawlak and Krawiec (2017a) is different in that their framework is a MILP. Similarly, other authors also leverage MILP (Schede et al., 2019), Integer Programming (IP) (Cozad et al., 2014), or LP (Sroka and Pawlak, 2018) frameworks for constraint learning. A number of authors (Pawlak and Krawiec, 2017b, 2018; Pawlak, 2019) also use a genetic programming framework to obtain constraints from data. In these different setups, the authors give the relative simplicity of their approaches as their key justification — no predictive models have to be learned, no constraints for feature extraction have to be included, and no complex embedding or reformulating procedures need to be followed. The disadvantage of these approaches however is that they don't make use of the learning power of predictive models which have been shown to be able to learn very complex decision boundaries well.

A key point to highlight is that the use of constraint learning in operations research is a relatively new development. Consequently, not many papers are seen to have a conceptually structured approach to doing optimization with constraint learning. Although several papers thoroughly describe their methodology of embedding a ML model into an optimization problem, they don't formalise the process of constraint learning from beginning to end. We believe that the formalization of this process is one of the important contributions of this paper to the literature.

4.2. Data Gathering and Preprocessing

Of the papers considered in this survey, the authors of (Thams et al., 2017) have the most detailed approach with regards to the data generation and preprocessing phase detailed in Section 3.2. Their procedure, also used in (Halilbašić et al., 2018; Venzke et al., 2020b), is designed to either make use of available data from sensors, or generate it using a simulator. Possible operating points of the system under consideration are fed into a simulator, and the output of the simulator is analysed and classified as either *stable* or *unstable*. These operating points along with their classifications are then stored in a database. While the data is only generated once, they attempt to keep the generated database as small as possible for computational reasons, while collecting enough data to train their predictive models. They do this

by focusing on generating points close to the boundary of their two-class problem, while neglecting operating points far away from the boundary. In this way they balance exploration and exploitation when generating data. In addition to this, class imbalance in the dataset is addressed using several strategies, and feature selection is done so that only easily accessible variables are considered.

With regards to following DoCE principles, the authors of (Lombardi et al., 2017) use factorial design to generate multiple training sets from a simulator. Data generation here, although time consuming, only needs to be done once. The data is also normalized so that the input features fall within a certain range. The approach developed in (Cozad et al., 2014) also takes DoCE principles into account by using either LHS or two-level factorial design. They also use adaptive sampling to iteratively refine the models generated and provide discussions on whether to use exploration-based or exploitation-based sampling. De Angelis et al. (2003) use small data collected using on-field survey (an expensive process) to find a probability distribution that is later used in a simulator that generates enough data to train the predictive model. In (Fahmi and Cremaschi, 2012) where several constraints are replaced using neural networks, the data is obtained from multiple simulators. Some transformation is done to the data to ensure that the distribution of the dependent variable is as close as possible to a uniform distribution. The authors here note that a large amount of data needed to train their neural networks for higher accuracy. Data may also be collated from multiple available sources as in the case of Bertsimas et al. (2016).

In terms of the sizes of the datasets surveyed, it can be seen that not many of the problems considered have large datasets. Data containing 500 or fewer examples are common. In (Pawlak and Krawiec, 2017a), the total amount of data generated using uniform sampling varied from 10 to 500 examples. Similarly, Schweidtmann and Mitsos (2019) use data with 500 or fewer examples. Data of a similar size is also seen in (Kudła and Pawlak, 2018), (Yang and Bequette, 2021) and several others. For one of the test problems in (Schweidtmann and Mitsos, 2019), only 46 examples are used. On the larger side, Chen et al. (2020) use data containing 10,000 instances, Venzke et al. (2020b) use 36,144 data points generated using LHS, while the data used in (Say et al., 2017) contains 100,000 data points. For certain constraint learning approaches, the availability of a large amount of data is a must. The local search approach used in (Sroka and Pawlak, 2018) is such an example, where the test set contains 500,000 examples generated via uniform

sampling.

Some approaches for constraint learning necessitate the significant preprocessing of the data. In order to learn ellipsoidal constraints in (Pawlak and Litwiniuk, 2021), the uniformly generated data first needs to be clustered and standardized, before Principal Component Analysis (PCA) is applied to the data. This clustering and PCA procedure shapes the data so that it can be represented using an ellipsoidal constraints. In other cases, although preprocessing is done, it is not crucial to the constraint learning approach (e.g. in (Biggs et al., 2018) and (Pawlak, 2019)).

To summarize, it is clear from the literature that with the exception of a few authors, not many papers have put much emphasis into the data generation and preprocessing phase. This is sometimes because the data is available, is neither big nor dirty, and has no need for preprocessing. For example, the situation of having high-dimensional data, where the number of examples is much less than the number of features, is only explicitly mentioned in (Mišić, 2020). In other cases however, the data step is an afterthought. Even when data is generated, this is cursorily mentioned without giving details or following exhaustive DoCE principles. We believe that our framework will help to provoke more thought on this phase of the constraint learning process.

4.3. Selection and Training of ML models

Before we begin this section, we would like to note that papers which use regression functions as the main constraint learning methods have been excluded from Table 1, as they are too numerous to be included in this paper. Regression with first- and second-order polynomials is used a lot in engineering, especially in combination with DoE and DoCE. We refer to Kleijnen (2015) for a detailed treatment. Kriging models are also used a lot in engineering, especially when the objective function values can be obtained by (expensive) computer simulations. The big advantage of kriging is that it has very high predictive power even for relatively small number of data points. The disadvantages of kriging are (i) it costs much computing time to obtain the kriging model (ii) the data points have to be space-filling (iii) embedding kriging functions in the optimization model makes it non-convex and hard to solve. See Forrester et al. (2008) for a detailed treatment on kriging.

It can be seen from Table 1 that ANNs and DTs are the most popularly used predictive models in constraint learning (with the exception of regression models). This could be because these methods have been shown to have high

Table 1: Methods used for constraint learning

	Neural Networks	Decision Trees	Random Forest	Other Ensemble	Support Vector Machines	Clustering	(M)ILP	Other
Bergman et al. (2019)	x							x
Biggs et al. (2018)		x	x					
Chen et al. (2020)	x							
Chi et al. (2007)					x			
Cozad et al. (2014)							x	x
Cremer et al. (2018)		x		x				
De Angelis et al. (2003)	x							
Fahmi and Cremaschi (2012)	x							x
Garg et al. (2018)					x			
Grimstad and Andersson (2019)	x							
Gutierrez-Martinez et al. (2010)	x							
Halilbašić et al. (2018)		x						
Jalali et al. (2019)					x			
Kudła and Pawlak (2018)		x						
Lombardi et al. (2017)	x	x						
Maragno et al. (2022)	x	x	x	x	x	x		x
Misić (2020)		x	x					
Paulus et al. (2021)	x							
Pawlak and Krawiec (2017a)							x	
Pawlak and Krawiec (2017b)								x
Pawlak and Krawiec (2018)								x
Pawlak (2019)								x
Pawlak and Litwiniuk (2021)								x
Pawlak and O'Neill (2021)						x		
Prat and Chatzivasileiadis (2020)		x						
Say et al. (2017)								
Schede et al. (2019)	x	x					x	
Schweidtmann and Mitsos (2019)	x							
Spyros (2020)	x	x						
Stoka and Pawlak (2018)								x
Thams et al. (2017)						x		
Venzke et al. (2020b)	x							
Verwer et al. (2017)		x						x
Xavier et al. (2021)					x			
Yang and Bequette (2021)	x							

representative power, can be used for both regression or classification, and generally perform well. Training, testing and evaluation is generally done as prescribed in the ML literature, and the methods are generally evaluated using metrics such as the Mean Square Error (MSE), accuracy, and so on. In terms of size of the ANNs, with the exception of Yang and Bequette (2021) who use a network with six hidden layers, most networks used are generally small. One reason for this could be to reduce the burden of embedding large networks as constraints. Consider the fully connected neural network shown in Figure 8. The parameter $w_{jk}^{[l]}$ represents the weight of input j to neuron k

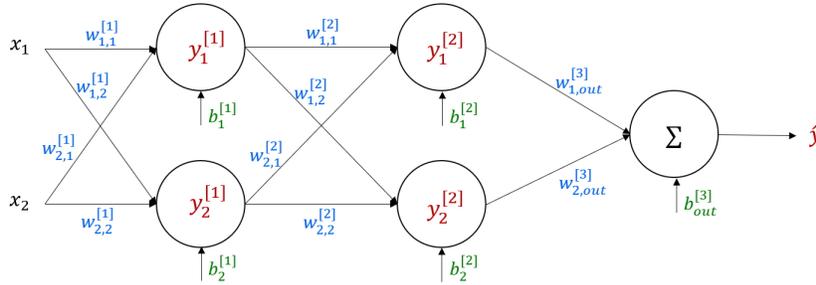


Figure 8: Example representation of a neural network

in layer l , while $b_j^{[l]}$ represents the bias of neuron j in layer l . The activation (or output) of a neuron j in layer l is given generally as

$$y_j^{[l]} = \sigma^{[l]} \left(\sum_{k=1}^{N_k} w_{jk}^{[l]} y_k^{[l-1]} + b_j^{[l]} \right), \quad (8)$$

where $\sigma^{[l]}$ is the activation function used in layer l and N_k is the number of neurons in layer $[l - 1]$.

The activation functions used in the hidden layers in the constraint learning literature are usually either the Rectified Linear Unit (ReLU) (e.g. in (Say et al., 2017)), the sigmoid function (e.g. in (Chen et al., 2020)), (some form of) the hyperbolic tangent function (e.g. in (Fahmi and Cremaschi, 2012)), or the softplus activation function (e.g. in (Chen et al., 2020)). These are shown in Figure 9. The output layer activation function is usually an identity linear function.

As the input of a current layer is the output of a previous one, using multiple layers results in nested activation functions in the final model. This could lead to issues regarding convexity and/or differentiability, therefore

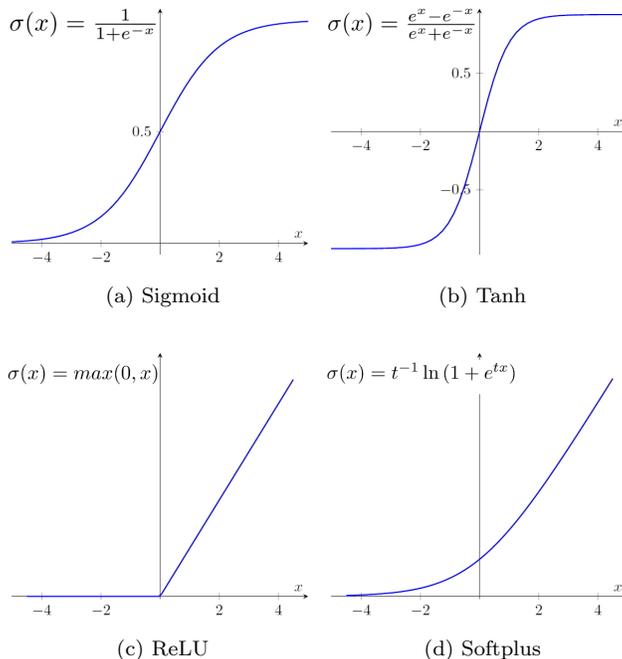


Figure 9: Commonly used activation functions

smaller networks are used. Another reason for small neural network sizes is that using larger networks are of no benefit in some cases. Lombardi et al. (2017) for example use a network with one hidden layer (with two neurons) as increasing the network size did not improve the solution quality of the optimization problem. To view the effect of network size on computation times, Schweidtmann and Mitsos (2019) varied both the number of neurons, and the number of layers. Results showed that the solution time increases approximately exponentially with the number of layers used to learn the constraints. Deep neural networks required more processing time than shallow neural networks for the same number of neurons.

An alternative to using small networks could be to use sparsification. Venzke et al. (2020b) use a network with three hidden layers and fifty neurons in each layer to learn constraints, however they reduce the size of the network by enforcing 70% of the weights to be zero. This sparsification procedure ensures that the network is not too large and also avoids over-fitting during training. Say et al. (2017) use a similar technique to remove connections to and from neurons with very small weights in order to strengthen the MILP

formulation. Another alternative could be to ensure that the constraint learning procedure results in a convex optimization problems. Chen et al. (2020) and Yang and Bequette (2021) use Input Convex Neural Networks (ICNN) (Amos et al., 2017) as their predictive models. An ICNN requires that all weights in the network are non-negative, and that all activation functions are convex and non-decreasing. Although these networks require more training iterations and have lower representation power than conventional networks, their use guarantees that the final problem is a convex optimization problem. This allows the authors to balance the trade-off between model accuracy and computational tractability.

In addition to ANNs, DTs are also popularly used in constraint learning. DTs are able to capture non-convex and discontinuous feasibility spaces, and provide these rules in an easily understandable format. Take for example, the decision tree in Figure 10 used to determine whether to go hiking or not.

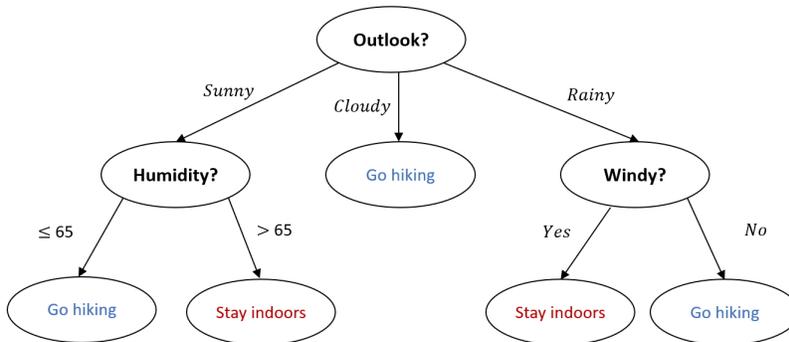


Figure 10: An example of a decision tree

It can easily be seen that if the outlook is sunny and the humidity is greater than 65, then we should stay indoors. For this reason, Thams et al. (2017) and Halilbašić et al. (2018) specifically choose decision trees as their method of choice to learn constraints in an optimal power flow problem. Their rationale is that as electricity is a crucial utility, using easily interpretable rules in their analysis can lower the barrier to power plant operators adopting their methods. This interpretability requirement could also favour the use of decision trees in the radiotherapy example of Section 2.2.

As random forests are a natural extension to using decision trees, they have also been used in constraint learning. Biggs et al. (2018) learned objective functions using random forest while Mišić (2020) optimized tree en-

sembles. AdaBoost (Freund and Schapire, 1997) (with DTs as base learners) is also used in (Cremer et al., 2018) to learn a probabilistic description of a constraint.

Support Vector Machines (SVM) can be expressed very clearly as optimization problems, and as a result, there is a significant amount of optimization literature on SVMs. Most of the literature however deal with subjects such as optimal feature selection (e.g. (Jiménez-Cordero et al., 2021)), or mathematical formulation and solution of the SVM problem (e.g. (Baldomero-Naranjo et al., 2020)), or improving predictions (e.g. (Yao et al., 2017)). There are however a few cases where SVM is used in the constraint learning process. Chi et al. (2007) uses SVM for regression with a 2nd order polynomial kernel function. The resulting learned constraint is of the form shown in (9) and is a linear combination of linear, quadratic and two-way interaction terms.

$$\begin{aligned}
 y = & w_0 + w_1x_1 + \dots + w_mx_m + w_{m+1}x_1^2 + \dots \\
 & + w_{2m}x_m^2 + w_{2m+1}x_1x_2 + w_{2m+2}x_1x_3 + \dots
 \end{aligned} \tag{9}$$

A feature selection process is further used to remove insignificant terms and improve model accuracy. Xavier et al. (2021) also use SVM with linear kernels to learn hyper-planes within which the solution of the optimization problem is very likely to exist.

Clustering is also used in constraint learning, although it is often used in addition to other methods. Pawlak and Litwiniuk (2021) combined clustering with PCA and ellipsoidal constraints to get an MIQCP, while Sroka and Pawlak (2018) combine clustering with local search to search the clusters for LP constraints that represent each cluster.

Combinations of predictive models are often used. Verwer et al. (2017) combine regression trees and regression models for learning revenue constraints for an auction optimization problem. In addition to using ANNs to learn most of the constraints required, Fahmi and Cremaschi (2012) also use non-linear regression models for specific constraints. Paulus et al. (2021) use both integer programming and neural networks to learn both the cost terms and the constraints for integer programs. Although Schede et al. (2019) use an LP formulation to learn polytopes that cover feasible examples, they also include a decision tree heuristic to model the importance of examples to be considered. Additionally, Maragno et al. (2022) are able to use different predictive models to learn different constraints of the same optimization problem in order to achieve a better performance.

There are also approaches which don't use predictive models to learn constraints. Pawlak and Krawiec (2017b, 2018) and Pawlak (2019) use genetic programming as their approach of choice, Pawlak and Krawiec (2017a) use a mixed-integer formulation and Cozad et al. (2014) combine simple basis functions using a MILP formulation to learn constraints.

Although RF is an ensemble method, the column "Other Ensemble" in the Table 1 is separated from the "Random Forest" column to highlight the fact that there is an opportunity to use other ensemble methods where different individual base learners can be used and their predictors combined. There is also the opportunity to try a wider variety of predictive models such as symbolic regression or Bayesian methods, especially if they are more easily represented using mathematical functions. Studies could be done to see if certain families of models perform better for certain types of problems. Finally, methods that provide some measure of uncertainty (such as confidence intervals or conditional probabilities) might also be chosen in order to be able to incorporate this uncertainty into the final mathematical model.

4.4. Resolution of the Optimization Model

With regression functions, the process of embedding them as constraints is relatively straightforward as the functions are simply a weighted sum of the decision variables (Bertsimas et al., 2016; Verwer et al., 2017). With other predictive models however, the embedding process is more involved. With ANNs for example, their embedding is also dependent on the activation functions used. In order to incorporate a trained neural network with ReLU activation functions as a constraint, the $max()$ function can be replaced by introducing binary variables b_j for each neuron in the network (e.g. in (Venzke et al., 2020b)). If the input to a neuron $\hat{y}_j > 0$, then $b_j = 1$ and (10a) and (10b) force the neuron output y_j to the input \hat{y}_j . On the other hand, if the input to the neuron $\hat{y}_j \leq 0$ then $b_j = 0$, and (10c) and (10d) force the output y_j to 0. The resulting MILP can then be solved with off-the-shelf solvers such as Gurobi or CPLEX.

$$y_j = \max(0, \hat{y}_j) \implies \begin{cases} y_j \leq \hat{y}_j - \hat{y}_j^{\min} (1 - b_j) & (10a) \\ y_j \geq \hat{y}_j & (10b) \\ y_j \leq \hat{y}_j^{\max} b_j & (10c) \\ y_j \geq 0 & (10d) \\ b_j \in \{0, 1\}^{N_j} & (10e) \end{cases}$$

The disadvantages of this approach are twofold. Firstly, large networks require a correspondingly large number of binary variables. Secondly, the bounds \hat{y}_j^{min} and \hat{y}_j^{max} have to be chosen to be as tight as possible for the MILP solver. Anderson et al. (2020) present a MIP formulation that produces strong relaxations for ReLUs. They start with the big-M formulation and then add strengthening inequalities as needed. Say et al. (2017) also encode ReLU into MILP using big-M constraints, but include reformulations to strengthen the bounds. They also use sparsification to remove neurons with very small weights, as constraints with very small coefficients can be difficult to solve for commercial solvers (D’andreagiovanni and Gleixner, 2016).

The embedding of trained ANNs as constraints is not always as complicated as mentioned above. Equation (8) can be used directly as a constraint, as long as the solver can handle the activation function σ . This approach is used in (Fahmi and Cremaschi, 2012) with the hyperbolic tangent activation function, as well as in (Gutierrez-Martinez et al., 2010) where the activation function is differentiable so that the constraint can be easily handled by conventional solvers. It is nevertheless noted in (Lombardi et al., 2017) that several commercial solvers rely on convexity for providing globally optimal results, and that this approach may result in locally optimum solutions depending on the σ used. In (Chen et al., 2020), the use of ICNN guarantees that the optimization problem is convex and will converge to a globally optimal solution. They do however speed up this process by using a smoothed version of the ReLU, called the Softplus activation function (see Figure 9d). A large value of the parameter t in the Softplus function allows it to be smooth enough, while still being convex. The final problem is solved using dual decomposition (Yu and Lui, 2006; Chiang et al., 2007). Yang and Bequette (2021) also use ICNNs to get convex optimization problems which are then solved using sequential quadratic programming.

Noticing the use of ANNs in various optimization applications, Schweidtmann and Mitsos (2019) provide an efficient method for deterministic global optimization problems which have ANNs embedded in them. As the only non-linearities of ANNs are their activation functions (the tanh function in their case), the tightness of their relaxations are influenced by the relaxations of the tanh functions. The key idea of their paper is to therefore hide the equations that describe the ANN and its variables from the Branch-and-Bound (B&B) solver using McCormick relaxations of the activation functions. These relaxations are once continuously differentiable, and the lower bounds are calculated by automatic propagation of McCormick relaxations through

the network equations. This “Reduced-Space (RS)” formulation results in significantly fewer variables and equalities than the original formulation (by two orders of magnitude in some cases), with significantly shorter computation times in most cases. They also see that shallow ANNs show much tighter relaxations than deep ANNs. Grimstad and Andersson (2019) also provide bound tightening procedures for ReLU networks to reduce solution times.

The embedding of kriging functions in the final model is done in a direct way: substituting the kriging function leads to an explicitly formulated, but highly non-convex constraint. First and second order derivatives can be calculated, and therefore standard nonlinear solvers can be used. However, due to the non-convexity, solvers might end up in local optima. See also den Hertog and Stehouwer (2002) and Stinstra and den Hertog (2008) for examples.

With DTs, the paths from root to leaf nodes can be represented using disjunctions of conjunctions. In Figure 10, there are three paths that lead us to go hiking which can be written as follows:

$$\begin{aligned} & (Outlook = Sunny \wedge Humidity \leq 65) \\ & \quad \vee (Outlook = Cloudy) \\ & \quad \vee (Outlook = Rainy \wedge Wind = No) \end{aligned} \tag{11}$$

These rules can either be linearized, or embedded as constraints using a framework like Generalized Disjunctive Programming (GDP) (Grossmann, 2009; GAMS Development Corporation, 2021). Although Lombardi et al. (2017) do not embed DTs into MINLPs due to the potential for poor bounds being obtained as a result of the extensive linearization of disjunctions required, other authors have. In (Kudła and Pawlak, 2018), branches of the tree from root to leaf nodes are encoded as linear constraints using a big-M formulation. Paths that end up in feasible decisions are denoted with the (+) sign (see Figure 11). These paths are then converted into linear constraints using auxiliary binary variables to show which path in the tree is followed. For example at the top of the tree, if a binary variable $b_1 = 1$, then c_1 is active and (12) holds. Otherwise if $b_1 = 0$, then c_2 is active and (13) holds.

$$c_1 : x_1 \leq a_1 \rightarrow x_1 \leq a_1 + M(1 - b_1) \tag{12}$$

$$c_2 : x_1 > a_1 \rightarrow x_1 > a_1 - Mb_1 \tag{13}$$

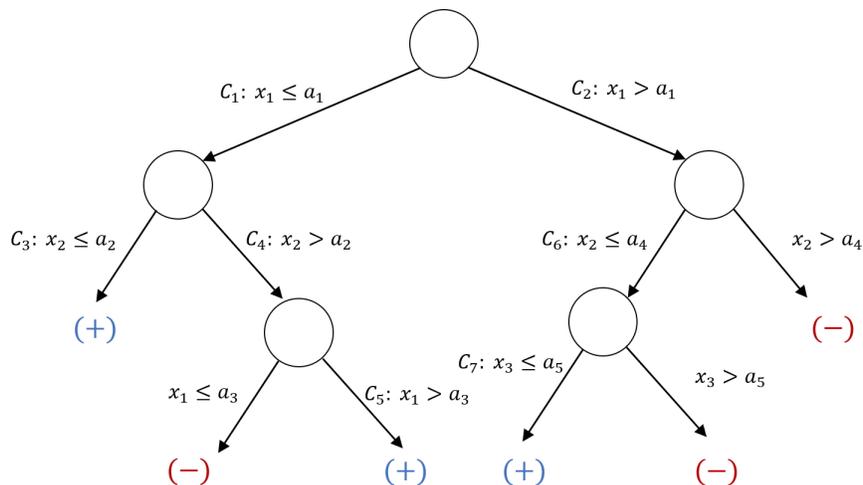


Figure 11: Extracting constraints from a decision tree (adapted from (Kudła and Pawlak, 2018)).

This process is applied to the whole tree until a set of linear constraints capturing the tree is obtained. A similar approach is used in (Verwer et al., 2017) to embed regression trees as constraints. In (Thams et al., 2017), each path from the root to a leaf node is represented by two constraints, with one binary variable per path to determine if a path is chosen. The resulting MILP is solved with Gurobi. Halilbašić et al. (2018) also follow a similar procedure except that the resulting problem is further reformulated as a Mixed-Integer Second Order Cone Program (MISOCP) because of the presence of additional non-linear constraints. The resulting MISOCP and MINLP are then solved using commercial solvers.

To embed random forests as constraints, multiple DTs can be used as above, with an additional constraint that selects a consensus based on the outcome chosen by most of the individual decision trees (Bonfietti et al., 2015). In (Biggs et al., 2018) the authors maximize a random forest as the objective. For each tree, they use logical constraints to determine which leaf a solution lies in. For the forest, they use a constraint to ensure that only one leaf for each tree is active. To ensure that the feasible set is bounded, and that the solution is not too dissimilar from the observed data, the solution is constrained to be within the convex hull. This is similar to the approach in (Mišić, 2020) who also have an objective function expressed as the prediction of a tree ensemble, and who also use some form of Benders reformulation with lazy constraint generation to solve the problem. Biggs et al. (2018) also

compare solving the MIP for the whole random forest, versus splitting the forest into a number of subsets, and solving each subset using the MIP, and a cross-validation heuristic procedure to achieve performance improvements. In (Cremer et al., 2018) where the AdaBoost ensemble learning method is used (with DTs as the weak learners), there is one set of constraints for each DT, and exactly one disjunction must be selected for each learner. Probability estimates for all learners are combined in a weighted sum, and the resulting MILP is solved using a commercial solver.

The advantage of using LP or MILP approaches to learn constraints is seen in this step of the framework, as there are no complex embedding procedures to be followed. In (Pawlak and Krawiec, 2017a) the resolution of the final optimization problem here is relatively easy as the entire constraint learning process was formulated as a MILP from the very beginning. This is also seen in papers like (Sroka and Pawlak, 2018) and (Schede et al., 2019). In (Pawlak and Litwiniuk, 2021), the data is first clustered, and then PCA is used to give each cluster an ellipsoidal shape. These ellipsoids can then be represented using quadratic constraints. Their approach produces Mixed Integer Quadratically Constrained Programs (MIQCP), which can then be solved by any solver that supports quadratic programming.

In summary, it can be seen that this step of the constraint learning process is the most complex one. Incorporating predictive models sometimes results in significant reformulation of the original model. Embedding of the predictive models should however be done in a way that takes full advantage of the optimization solver’s capabilities. This can be done by providing gradients or other useful information directly to the solver, or reformulating the problem to become linear, convex, conic quadratic, etc. Constraint learning approaches which use MILP or LP frameworks such as in (Pawlak and Krawiec, 2017a) avoid these complex embedding procedures, however the disadvantages of this approach could include a lack of flexibility and limited learning ability. There is therefore a trade-off between simplicity and performance to be considered. Wherever possible, it is desirable to keep the final optimization problem in the same complexity class as the original problem.

4.5. Verification and Improvement of the Optimization Model

The validation of the final optimization model is straightforward when dealing with benchmark problems, as the ground-truth is already known. This is seen in a lot of the literature. In (Pawlak and Krawiec, 2017a, 2018)

and (Pawlak, 2019) for example, the use of benchmark problems allows for the comparison of the number of synthesized constraints with the actual number of constraints from the benchmark problems. They are also able to assess the similarity of the syntaxes by computing the mean angle between weight vectors of the corresponding constraints in the synthesized and actual models. The feasible regions of the synthesized models can also be compared to those of the actual feasible regions using the Jaccard Index (Jaccard, 1912), as it is used to measure the similarity between sets. This is used in (Kudła and Pawlak, 2018; Pawlak and Krawiec, 2017a) and others. Visual comparisons (limited to 3-dimensions) have also been used to illustrate how well a constraint learning approach captures the real feasible region (e.g. (Pawlak and Litwiniuk, 2021)). In (Schede et al., 2019), the authors compute the probabilities that feasible (and infeasible) points of the benchmark problem lie in the feasible (and infeasible) region of the learned problem. They also compare the differences between the optimum values of the original and learned problems. Gutierrez-Martinez et al. (2010); Cremer et al. (2018) and Schweidtmann and Mitsos (2019) similarly also test their approaches on benchmark problems, and compare their results to the known solutions of the problems. In addition to using benchmark problems, Xavier et al. (2021) also evaluate the performance of the method by using out-of-distribution data to measure the robustness against moderate dataset shift.

Using synthetic or well-known benchmark problems allows one to accurately determine how far learned constraints are from the real ones, however this is not always possible. Consequently, other model verification procedures are seen in the literature. Simulators can also be used to evaluate constraint learning approaches by comparing learned solutions to solutions from the simulator (e.g. in (Verwer et al., 2017)).

Authors also compare the constraint learning approaches with the typical approaches used to solve those types of problems. This is seen in (Say et al., 2017), (Thams et al., 2017), (Chen et al., 2020) among others. Cozad et al. (2014) implement model improvement by iteratively generating new data and updating the model with these data. The process continues until the error is less than a specified tolerance value. An overview of the approaches used for verification and improvement of the learned models is given in Table 2.

Table 2: Approaches used for model verification and improvement

	Benchmark Problems	Compare Constraints and/or Feasible Regions	Simulator Results	Compare to Other Approaches	Adaptive Sampling	Other
Bergman et al. (2019)	x					
Chen et al. (2020)	x		x	x		
Chi et al. (2007)			x			
Cozad et al. (2014)			x		x	
Cremer et al. (2018)	x		x	x		
De Angelis et al. (2003)			x		x	
Fahmi and Cremaschi (2012)			x			
Garg et al. (2018)	x			x		
Grimstad and Andersson (2019)	x					
Gutierrez-Martinez et al. (2010)	x			x		
Halilbašić et al. (2018)	x			x		
Jalali et al. (2019)	x			x		
Kudła and Pawlak (2018)	x	x		x		
Lombardi et al. (2017)			x	x		
Mišić (2020)				x		
Paulus et al. (2021)	x					
Pawlak and Krawiec (2017a)	x	x				
Pawlak and Krawiec (2017b)	x	x				
Pawlak and Krawiec (2018)	x	x				
Pawlak (2019)	x	x				
Pawlak and Litwiniuk (2021)	x	x		x		
Pawlak and O'Neill (2021)	x	x		x		
Prat and Chatzivasileiadis (2020)	x			x		
Say et al. (2017)	x			x		
Schede et al. (2019)	x	x		x		
Schweidtmann and Mitsos (2019)	x			x		
Spyros (2020)	x			x		
Stoka and Pawlak (2018)	x	x		x		
Thams et al. (2017)	x			x		
Venzke et al. (2020b)	x			x		
Verwer et al. (2017)						x
Xavier et al. (2021)	x					
Yang and Bequette (2021)	x			x		

5. Opportunities for Further Research

As the process of constraint learning is relatively new, it stands to reason that there will be several areas where further research can be carried out. Following, we summarize some of these opportunities for further research in light of the steps of our proposed framework.

Data Gathering and Preprocessing: As seen in Section 4.2, the data gathering and preprocessing phase of the constraint learning process is often an afterthought in optimization. There is therefore an opportunity to understand the effects of the data generation and preprocessing phase on the solution of the optimization problem. Is there a particular strategy (e.g. DoCE, Kriging, etc.) that performs best for constraint learning? It is also necessary to understand the size and type of data that is needed in order to get a good predictive model for constraint learning, and not just for prediction.

Selection and Training: Firstly, the choice of model to use to learn constraints should also be further investigated. Most of the approaches used to learn constraints are the most commonly known methods (regression, ANNs, DTs, etc.), however there is an opportunity to see how less well-known approaches perform in terms of solution quality, as well as ease of embedding as constraints. Secondly, predictive models often provide a measure of the uncertainty of their predictions. Neural Network classifiers, for example, can have Softmax scores while Decision Trees report the number of misclassified examples at their leaf nodes. Although Cremer et al. (2018) have used probabilistic information from their classifier to good results, further research on how to incorporate these measures of uncertainty into optimization problems should be carried out. Approaches such as robust optimization or stochastic programming might be useful here. Thirdly, only a few approaches have used ensemble methods to learn constraints. While not all problems will require the additional predictive power that an ensemble approach brings, it is to be expected that as the best predictive models in ML applications are often ensembles, it makes sense to incorporate their superior abilities into constraint learning. Finally, in Section 3.1, we make the distinction between constraint learning and constraint function learning. Further research is required on which approach performs best, either generally or on a case-by-case basis. The relationship between the two approaches should be investigated. When do two methods result in the same optimal solution? In other words, when

do they define the same feasible region?

Resolution of the Optimization Model: Several approaches for embedding ML models as constraints have been seen, however can these models be embedded in more efficient ways? The difficulty of using larger neural networks or larger decision trees to learn constraints might be overcome if more efficient embedding procedures are developed. In (Schweidtmann and Mitsos, 2019), difficult parts of the neural network (tanh activation functions) were hidden from the B&B solver. Further research on hiding unnecessary information from solvers is needed. Further research is also needed on providing additional useful information (such as pre-computed derivatives) to commercial solvers so as to improve solution time and quality. This will be useful when the incorporation of predictive models make these computations impractical or expensive.

Verification and Improvement: In the literature reviewed, methods for model verification and improvement, usually consist of either comparing to benchmark problems, or to other competing approaches. It could be promising to look into developing formal approaches for verifying learned constraint models. A framework for verifying neural network behavior was developed in (Venzke and Chatzivasileiadis, 2020), and research could be done to see if similar frameworks can be applied during the constraint learning process. In terms of model improvement, more focus should be put on improving models via processes such as active learning or adaptive sampling. In such cases, it would be interesting to see the effects (if any) of exploration versus exploitation when generating data on the overall solution of the optimization problem. Questions also arise with respect to the robustness of the solution – What is the sensitivity of the optimal solution with respect to the uncertainty in the learned constraint? How is the robustness of the solution related to the noise in the dataset? A structured and formal approach for verifying optimization problems with learned constraints can help to answer questions such as these.

6. Conclusions

Reading any of the literature surveyed in this paper will clearly show the benefits of OCL. OCL helps to capture constraints which would otherwise have been difficult to capture, providing there is data available. These ben-

efits are clearly stated by Halilbašić et al. (2018) where their data-driven approach increased the feasible space, identified the global optimum and was solved in less time than conventional methods. There is however a need to go about the process of OCL in an organized manner, in order to avoid potential pitfalls. We have therefore provided a framework for OCL, which we believe will help to formalize and direct the process of OCL. We start with the setup of the conceptual model, and then look at the process of collecting and pre-processing the data. We also look at the factors that influence the methods and algorithms used to learn the constraints, and the techniques by which these learned constraints are embedded into the optimization models. Finally, we highlight the steps for verification and improvement of these models. We also provide two running examples — the WFP diet palatability problem, and the radiotherapy optimization problem — and used these examples to illustrate our framework. Besides providing the framework, we have reviewed the literature in light of the different steps of the framework, highlighting common trends in constraint learning, as well as possible areas for future research. It is our belief that this paper will help to guide future efforts in OCL, and be of benefit to the wider OR community.

Acknowledgements

This work was supported by the Dutch Scientific Council (NWO) grant OCENW.GROOT.2019.015, Optimization for and with Machine Learning (OPTIMAL). The authors would also like to thank Ilker Birbil for his comments during the preparation of the manuscript.

References

- Abbasi, B., Babaei, T., Hosseinifard, Z., Smith-Miles, K., Dehghani, M., 2020. Predicting solutions of large-scale optimization problems via machine learning: A case study in blood supply chain management. *Computers & Operations Research* 119, 104941.
- Amos, B., Xu, L., Kolter, J.Z., 2017. Input convex neural networks, in: *International Conference on Machine Learning*, PMLR. pp. 146–155.
- Anderson, R., Huchette, J., Ma, W., Tjandraatmadja, C., Vielma, J.P., 2020. Strong mixed-integer programming formulations for trained neural networks. *Mathematical Programming* 183, 3–39.

- Bagloee, S.A., Asadi, M., Sarvi, M., Patriksson, M., 2018. A hybrid machine-learning and optimization method to solve bi-level problems. *Expert Systems with Applications* 95, 142–152.
- Baldomero-Naranjo, M., Martínez-Merino, L.I., Rodríguez-Chía, A.M., 2020. Tightening big Ms in integer programming formulations for support vector machines with ramp loss. *European Journal of Operational Research* 286, 84–100.
- Beldiceanu, N., Simonis, H., 2016. Modelseeker: Extracting global constraint models from positive examples, in: *Data Mining and Constraint Programming*. Springer, pp. 77–95.
- Bengio, Y., Lodi, A., Prouvost, A., 2021. Machine learning for combinatorial optimization: a methodological tour d’horizon. *European Journal of Operational Research* 290, 405–421.
- Bennett, K.P., Parrado-Hernández, E., 2006. The interplay of optimization and machine learning research. *Journal of Machine Learning Research* 7, 1265–1281.
- Bergman, D., Huang, T., Brooks, P., Lodi, A., Raghunathan, A.U., 2019. Janos: An integrated predictive and prescriptive modeling framework. [arXiv:1911.09461](https://arxiv.org/abs/1911.09461).
- Bergstra, J., Bardenet, R., Bengio, Y., Kégl, B., 2011. Algorithms for hyper-parameter optimization, in: *Shawe-Taylor, J., Zemel, R., Bartlett, P., Pereira, F., Weinberger, K.Q. (Eds.), Advances in Neural Information Processing Systems*, Curran Associates, Inc.. pp. 2546–2554. URL: <https://bit.ly/3pmrL6e>.
- Bertsimas, D., Dunn, J.W., 2019. Machine learning under a modern optimization lens. *Dynamic Ideas LLC*.
- Bertsimas, D., Kallus, N., 2020. From predictive to prescriptive analytics. *Management Science* 66, 1025–1044.
- Bertsimas, D., O’Hair, A., Relyea, S., Silberholz, J., 2016. An analytics approach to designing combination chemotherapy regimens for cancer. *Management Science* 62, 1511–1531.

- Bessiere, C., Coletta, R., Freuder, E.C., O’Sullivan, B., 2004. Leveraging the learning power of examples in automated constraint acquisition, in: International Conference on Principles and Practice of Constraint Programming, Springer. pp. 123–137.
- Biggs, M., Hariss, R., Perakis, G., 2018. Optimizing objective functions determined from random forests. Available at SSRN 2986630 .
- Bonfietti, A., Lombardi, M., Milano, M., 2015. Embedding decision trees and random forests in constraint programming, in: International Conference on AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems, Springer. pp. 74–90.
- Breiman, L., Friedman, J., Stone, C.J., Olshen, R.A., 1984. Classification and regression trees. CRC press.
- Chen, Y., Shi, Y., Zhang, B., 2020. Input convex neural networks for optimal voltage regulation. arXiv preprint arXiv:2002.08684 .
- Chi, H.M., Ersoy, O.K., Moskowitz, H., Ward, J., 2007. Modeling and optimizing a vendor managed replenishment system using machine learning and genetic algorithms. European Journal of Operational Research 180, 174–193.
- Chiang, M., Low, S.H., Calderbank, A.R., Doyle, J.C., 2007. Layering as optimization decomposition: A mathematical theory of network architectures. Proceedings of the IEEE 95, 255–312.
- Conn, A., Scheinberg, K., Vicente, L., 2009. Introduction to derivative-free optimization, in: MPS-SIAM Series on Optimization.
- Cozad, A., Sahinidis, N.V., Miller, D.C., 2014. Learning surrogate models for simulation-based optimization. AIChE Journal 60, 2211–2227.
- Cremer, J.L., Konstantelos, I., Tindemans, S.H., Strbac, G., 2018. Data-driven power system operation: Exploring the balance between cost and risk. IEEE Transactions on Power Systems 34, 791–801.
- Crombecq, K., Laermans, E., Dhaene, T., 2011. Efficient space-filling and non-collapsing sequential design strategies for simulation-based modeling. European Journal of Operational Research 214, 683–696. URL: <https://www.sciencedirect.com/science/article/pii/>

S0377221711004577, doi:<https://doi.org/10.1016/j.ejor.2011.05.032>.

- Dantzig, G.B., 1990. The diet problem. *Interfaces* 20, 43–47.
- De Angelis, V., Felici, G., Impelluso, P., 2003. Integrating simulation and optimisation in health care centre management. *European Journal of Operational Research* 150, 101–114.
- Demirovic, E., Stuckey, P.J., Bailey, J., Chan, J., Leckie, C., Ramamohanarao, K., Guns, T., 2019. Predict+ optimise with ranking objectives: Exhaustively learning linear functions., in: *International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 1078–1085.
- den Hertog, D., Stehouwer, H., 2002. Optimizing color picture tubes by high-cost non-linear programming. *European Journal of Operational Research* 140, 197–211. Pagination: 14.
- Deng, Y., Liu, J., Sen, S., 2018. Coalescing data and decision sciences for analytics, in: *Recent Advances in Optimization and Modeling of Contemporary Problems*. INFORMS, pp. 20–49.
- Donti, P., Amos, B., Kolter, J.Z., 2017. Task-based end-to-end model learning in stochastic optimization, in: *Advances in Neural Information Processing Systems*, pp. 5484–5494.
- D’andreagiovanni, F., Gleixner, A.M., 2016. Towards an accurate solution of wireless network design problems, in: *International Symposium on Combinatorial Optimization*, Springer. pp. 135–147.
- Eikelder, S., Ajdari, A., Bortfeld, T., Hertog, D.d., et al., 2019. Adjustable robust treatment-length optimization in radiation therapy. *arXiv preprint arXiv:1906.12116* .
- Elmachtoub, A.N., Grigas, P., 0. Smart “predict, then optimize”. *Management Science* 0, null. URL: <https://doi.org/10.1287/mnsc.2020.3922>, doi:10.1287/mnsc.2020.3922, arXiv:<https://doi.org/10.1287/mnsc.2020.3922>.
- Fahmi, I., Cremaschi, S., 2012. Process synthesis of biodiesel production plant using artificial neural networks as the surrogate models. *Computers & Chemical Engineering* 46, 105–123.

- Fischetti, M., Fraccaro, M., 2019. Machine learning meets mathematical optimization to predict the optimal production of offshore wind parks. *Computers & Operations Research* 106, 289–297.
- Forrester, A., Sobester, A., Keane, A., 2008. *Engineering design via surrogate modelling: a practical guide*. Wiley.
- Freund, Y., Schapire, R.E., 1997. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences* 55, 119–139.
- Galassi, A., Lombardi, M., Mello, P., Milano, M., 2018. Model agnostic solution of csps via deep learning: A preliminary study, in: *International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, Springer. pp. 254–262.
- GAMS Development Corporation, 2021. Disjunctive programming. URL: https://www.gams.com/latest/docs/UG_EMP_DisjunctiveProgramming.html.
- Garg, A., Jalali, M., Kekatos, V., Gatsis, N., 2018. Kernel-based learning for smart inverter control, in: *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, IEEE. pp. 875–879.
- Garud, S.S., Karimi, I.A., Kraft, M., 2017. Design of computer experiments: A review. *Computers & Chemical Engineering* 106, 71–95. URL: <https://www.sciencedirect.com/science/article/pii/S0098135417302090>, doi:<https://doi.org/10.1016/j.compchemeng.2017.05.010>. eSCAPE-26.
- Gay, H.A., Niemierko, A., 2007. A free program for calculating eud-based ntcp and tcp in external beam radiotherapy. *Physica Medica* 23, 115–125. URL: <https://www.sciencedirect.com/science/article/pii/S1120179707000609>, doi:<https://doi.org/10.1016/j.ejmp.2007.07.001>.
- Gilan, S.S., Dilkina, B., 2015. Sustainable building design: A challenge at the intersection of machine learning and design optimization, in: *AAAI Workshop: Computational Sustainability*, pp. 101–106.

- Grimstad, B., Andersson, H., 2019. Relu networks as surrogate models in mixed-integer linear programs. *Computers & Chemical Engineering* 131, 106580.
- Grossmann, I.E., 2009. *Generalized disjunctive programming*. Springer US, Boston, MA. pp. 1180–1183. doi:10.1007/978-0-387-74759-0_203.
- Gutierrez-Martinez, V.J., Cañizares, C.A., Fuerte-Esquivel, C.R., Pizano-Martinez, A., Gu, X., 2010. Neural-network security-boundary constrained optimal power flow. *IEEE Transactions on Power Systems* 26, 63–72.
- Halilbašić, L., Thams, F., Venzke, A., Chatzivasileiadis, S., Pinson, P., 2018. Data-driven security-constrained ac-opf for operations and markets, in: 2018 Power Systems Computation Conference (PSCC), IEEE. pp. 1–7.
- Han, B., Shang, C., Huang, D., 2021. Multiple kernel learning-aided robust optimization: Learning algorithm, computational tractability, and usage in multi-stage decision-making. *European Journal of Operational Research* 292, 1004–1018.
- Hoffmann, A.L., den Hertog, D., Siem, A.Y., Kaanders, J.H., Huizenga, H., 2008. Convex reformulation of biologically-based multi-criteria intensity-modulated radiation therapy optimization including fractionation effects. *Physics in Medicine & Biology* 53, 6345.
- Hottung, A., Tanaka, S., Tierney, K., 2020. Deep learning assisted heuristic tree search for the container pre-marshalling problem. *Computers & Operations Research* 113, 104781.
- Jaccard, P., 1912. The distribution of the flora in the alpine zone. 1. *New Phytologist* 11, 37–50.
- Jain, P., Kar, P., 2017. Non-convex optimization for machine learning. *Foundations and Trends in Machine Learning* 10, 142–336. URL: <http://dx.doi.org/10.1561/22000000058>, doi:10.1561/22000000058.
- Jalali, M., Kekatos, V., Gatsis, N., Deka, D., 2019. Designing reactive power control rules for smart inverters using support vector machines. *IEEE Transactions on Smart Grid* 11, 1759–1770.

- Jiménez-Cordero, A., Morales, J.M., Pineda, S., 2021. A novel embedded min-max approach for feature selection in nonlinear support vector machine classification. *European Journal of Operational Research* 293, 24–35.
- Khalil, E., Le Bodic, P., Song, L., Nemhauser, G., Dilkina, B., 2016. Learning to branch in mixed integer programming. *Proceedings of the AAAI Conference on Artificial Intelligence* 30, 1. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/10080>.
- Kim, J.H., 2008. Linac-based high-precision radiotherapy: Radiosurgery, image-guided radiotherapy, and respiratory-gated radiotherapy. *Journal of the Korean Medical Association* 51, 612–618. doi:10.5124/jkma.2008.51.7.612.
- Kleijnen, J., 2015. *Design and analysis of simulation experiments*. Springer.
- Kolb, S.M., 2016. Learning constraints and optimization criteria, in: *Workshops at the Thirtieth AAAI Conference on Artificial Intelligence*, pp. 403–409.
- Kotary, J., Fioretto, F., Van Hentenryck, P., Wilder, B., 2021. End-to-end constrained optimization learning: A survey. *arXiv preprint arXiv:2103.16378* .
- Kudła, P., Pawlak, T.P., 2018. One-class synthesis of constraints for mixed-integer linear programming with c4.5 decision trees. *Applied Soft Computing* 68, 1–12.
- Lombardi, M., Milano, M., 2018. Boosting combinatorial problem modeling with machine learning, in: *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18, International Joint Conferences on Artificial Intelligence Organization*. pp. 5472–5478. URL: <https://doi.org/10.24963/ijcai.2018/772>, doi:10.24963/ijcai.2018/772.
- Lombardi, M., Milano, M., Bartolini, A., 2017. Empirical decision model learning. *Artificial Intelligence* 244, 343–367.
- Maragno, D., Wiberg, H., Bertsimas, D., Birbil, I., den Hertog, D., Fajemisin, A.O., 2022. *Optimization with constraint learning*. preprint .

- Matheron, G., 1963. Principles of geostatistics. *Economic Geology* 58, 1246–1266.
- Mišić, V.V., 2020. Optimization of tree ensembles. *Operations Research* 68, 1605–1624.
- Müller, J., 2009. Automation and optimization of fluence-modulated radiation therapy (IMRT) for head and neck cancer. Ph.D. thesis. Saarland University, Germany.
- Nagata, Y., Chu, K.H., 2003. Optimization of a fermentation medium using neural networks and genetic algorithms. *Biotechnology Letters* 25, 1837–1842.
- Nascimento, C.A.O., Giudici, R., Guardani, R., 2000. Neural network based approach for optimization of industrial chemical processes. *Computers & Chemical Engineering* 24, 2303–2314.
- Ng, A., 2020. Neural networks and deep learning. URL: <https://www.coursera.org/specializations/deep-learning>.
- O’Sullivan, B., 2010. Automated modelling and solving in constraint programming, in: *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence AAAI-10*, pp. 1493–1497.
- Padmanabhan, B., Tuzhilin, A., 2003. On the use of optimization for data mining: Theoretical interactions and ecrm opportunities. *Management Science* 49, 1327–1343.
- Paulus, A., Rolínek, M., Musil, V., Amos, B., Martius, G., 2021. Comboptnet: Fit the right np-hard problem by learning integer programming constraints. arXiv preprint arXiv:2105.02343 .
- Pawlak, T.P., 2019. Synthesis of mathematical programming models with one-class evolutionary strategies. *Swarm and Evolutionary Computation* 44, 335–348.
- Pawlak, T.P., Krawiec, K., 2017a. Automatic synthesis of constraints from examples using mixed integer linear programming. *European Journal of Operational Research* 261, 1141–1157.

- Pawlak, T.P., Krawiec, K., 2017b. Synthesis of mathematical programming constraints with genetic programming, in: European Conference on Genetic Programming, Springer. pp. 178–193.
- Pawlak, T.P., Krawiec, K., 2018. Synthesis of constraints for mathematical programming with one-class genetic programming. *IEEE Transactions on Evolutionary Computation* 23, 117–129.
- Pawlak, T.P., Litwiniuk, B., 2021. Ellipsoidal one-class constraint acquisition for quadratically constrained programming. *European Journal of Operational Research* 293, 36–49.
- Pawlak, T.P., O’Neill, M., 2021. Grammatical evolution for constraint synthesis for mixed-integer linear programming. *Swarm and Evolutionary Computation* 64, 100896.
- 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.
- Peters, K., Silva, S., Gonçalves, R., Kavelj, M., Fleuren, H., den Hertog, D., Ergun, O., Freeman, M., 2021. The nutritious supply chain: Optimizing humanitarian food assistance. *Inform Journal on Optimization* 3, 200–226.
- Prat, E., Chatzivasileiadis, S., 2020. Learning active constraints to efficiently solve bilevel problems. *arXiv preprint arXiv:2010.06344* .
- Santner, T.J., Williams, B.J., Notz, W.I., Williams, B.J., 2003. The design and analysis of computer experiments. volume 1. Springer.
- Say, B., Devriendt, J., Nordström, J., Stuckey, P.J., 2020. Theoretical and experimental results for planning with learned binarized neural network transition models, in: International Conference on Principles and Practice of Constraint Programming, Springer. pp. 917–934.
- Say, B., Sanner, S., 2018. Planning in factored state and action spaces with learned binarized neural network transition models., in: Proceedings of the

Twenty-Seventh International Joint Conference on Artificial Intelligence IJCAI-18, pp. 4815–4821.

Say, B., Wu, G., Zhou, Y.Q., Sanner, S., 2017. Nonlinear hybrid planning with deep net learned transition models and mixed-integer linear programming, in: Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence IJCAI-17, pp. 750–756.

Schede, E.A., Kolb, S., Teso, S., 2019. Learning linear programs from data, in: 2019 IEEE 31st International Conference on Tools with Artificial Intelligence (ICTAI), IEEE. pp. 1019–1026.

Schweidtmann, A.M., Mitsos, A., 2019. Deterministic global optimization with artificial neural networks embedded. *Journal of Optimization Theory and Applications* 180, 925–948.

Scott, J.G., Sedor, G., Scarborough, J.A., Kattan, M.W., Peacock, J., Grass, G.D., Mellon, E.A., Thapa, R., Schell, M., Waller, A., et al., 2021. Personalizing radiotherapy prescription dose using genomic markers of radiosensitivity and normal tissue toxicity in nscl. *Journal of Thoracic Oncology* 16, 428–438.

Shang, C., Huang, X., You, F., 2017. Data-driven robust optimization based on kernel learning. *Computers & Chemical Engineering* 106, 464–479.

Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al., 2016. Mastering the game of go with deep neural networks and tree search. *Nature* 529, 484–489.

Spyros, C., 2020. From decision trees and neural networks to milp: Power system optimization considering dynamic stability constraints, in: 2020 European Control Conference (ECC), IEEE. pp. 594–594.

Sroka, D., Pawlak, T.P., 2018. One-class constraint acquisition with local search, in: Proceedings of the Genetic and Evolutionary Computation Conference, pp. 363–370.

Stigler, G.J., 1945. The cost of subsistence. *Journal of Farm Economics* 27, 303–314.

- Stinstra, E., den Hertog, D., 2008. Robust optimization using computer experiments. *European Journal of Operational Research* 191, 816–837. URL: <https://www.sciencedirect.com/science/article/pii/S0377221707004316>, doi:<https://doi.org/10.1016/j.ejor.2007.03.048>.
- Sun, S., Cao, Z., Zhu, H., Zhao, J., 2019. A survey of optimization methods from a machine learning perspective. *arXiv:1906.06821*.
- Thams, F., Halilbasic, L., Pinson, P., Chatzivasileiadis, S., Eriksson, R., 2017. Data-driven security-constrained opf, in: *Proc. 10th Bulk Power Syst. Dyn. Control Symp.*, pp. 1–10.
- Tucker, S.L., Li, M., Xu, T., Gomez, D., Yuan, X., Yu, J., Liu, Z., Yin, M., Guan, X., Wang, L.E., Wei, Q., Mohan, R., Vinogradskiy, Y., Martel, M., Liao, Z., 2013. Incorporating single-nucleotide polymorphisms into the lyman model to improve prediction of radiation pneumonitis. *International Journal of Radiation Oncology*Biography*Physics* 85, 251–257. URL: <http://www.sciencedirect.com/science/article/pii/S0360301612002337>, doi:<https://doi.org/10.1016/j.ijrobp.2012.02.021>.
- UNWFP, U., 2021. The wfp food basket. URL: <https://www.wfp.org/wfp-food-basket>.
- Vaclavik, R., Novak, A., Sucha, P., Hanzlek, Z., 2018. Accelerating the branch-and-price algorithm using machine learning. *European Journal of Operational Research* 271, 1055 – 1069.
- Venzke, A., Chatzivasileiadis, S., 2020. Verification of neural network behaviour: Formal guarantees for power system applications. *IEEE Transactions on Smart Grid* 12, 383–397.
- Venzke, A., Qu, G., Low, S., Chatzivasileiadis, S., 2020a. Learning optimal power flow: Worst-case guarantees for neural networks, in: *2020 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, IEEE. pp. 1–7.
- Venzke, A., Viola, D.T., Mermet-Guyennet, J., Misyris, G.S., Chatzivasileiadis, S., 2020b. Neural networks for encoding dynamic security-

- constrained optimal power flow to mixed-integer linear programs. arXiv preprint arXiv:2003.07939 .
- Verwer, S., Zhang, Y., Ye, Q.C., 2017. Auction optimization using regression trees and linear models as integer programs. *Artificial Intelligence* 244, 368–395.
- Villarrubia, G., De Paz, J.F., Chamoso, P., De la Prieta, F., 2018. Artificial neural networks used in optimization problems. *Neurocomputing* 272, 10–16.
- Xavier, Á.S., Qiu, F., Ahmed, S., 2021. Learning to solve large-scale security-constrained unit commitment problems. *INFORMS Journal on Computing* 33, 739–756.
- Xu, S., Liu, H., Wang, X., Jiang, X., 2014. A robust error-pursuing sequential sampling approach for global metamodeling based on voronoi diagram and cross validation. *Journal of Mechanical Design* 136, 071009.
- Yang, L., Shami, A., 2020. On hyperparameter optimization of machine learning algorithms: Theory and practice. *Neurocomputing* 415, 295–316. URL: <http://www.sciencedirect.com/science/article/pii/S0925231220311693>, doi:<https://doi.org/10.1016/j.neucom.2020.07.061>.
- Yang, S., Bequette, B.W., 2021. Optimization-based control using input convex neural networks. *Computers & Chemical Engineering* 144, 107143.
- Yao, X., Crook, J., Andreeva, G., 2017. Enhancing two-stage modelling methodology for loss given default with support vector machines. *European Journal of Operational Research* 263, 679–689. URL: <https://www.sciencedirect.com/science/article/pii/S0377221717304459>, doi:<https://doi.org/10.1016/j.ejor.2017.05.017>.
- Yu, W., Lui, R., 2006. Dual methods for nonconvex spectrum optimization of multicarrier systems. *IEEE Transactions on Communications* 54, 1310–1322.