

MARK TURNER<sup>1</sup>, ANTONIA CHMIELA<sup>2</sup>, THORSTEN KOCH<sup>3</sup>,  
MICHAEL WINKLER

# PySCIPOpt-ML: Embedding Trained Machine Learning Models into Mixed-Integer Programs

---

<sup>1</sup>  [0000-0001-7270-1496](https://orcid.org/0000-0001-7270-1496)

<sup>2</sup>  [0000-0002-4809-2958](https://orcid.org/0000-0002-4809-2958)

<sup>3</sup>  [0000-0002-1967-0077](https://orcid.org/0000-0002-1967-0077)

Zuse Institute Berlin  
Takustr. 7  
14195 Berlin  
Germany

Telephone: +49 30 84185-0  
Telefax: +49 30 84185-125

E-mail: [bibliothek@zib.de](mailto:bibliothek@zib.de)  
URL: <http://www.zib.de>


ZIB-Report (Print) ISSN 1438-0064  
ZIB-Report (Internet) ISSN 2192-7782


---

# PYSCIPOPT-ML: EMBEDDING TRAINED MACHINE LEARNING MODELS INTO MIXED-INTEGER PROGRAMS

---

A PREPRINT

 **Mark Turner\***  
turner@zib.de

 **Antonia Chmiela\***  
chmiela@zib.de

 **Thorsten Koch<sup>†</sup>\***  
koch@zib.de

**Michael Winkler<sup>‡</sup>**  
winkler@gurobi.com

December 20, 2023

## ABSTRACT

A standard tool for modelling real-world optimisation problems is mixed-integer programming (MIP). However, for many of these problems there is either incomplete information describing variable relations, or the relations between variables are highly complex. To overcome both these hurdles, machine learning (ML) models are often used and embedded in the MIP as surrogate models to represent these relations. Due to the large amount of available ML frameworks, formulating ML models into MIPs is highly non-trivial. In this paper we propose a tool for the automatic MIP formulation of trained ML models, allowing easy integration of ML constraints into MIPs. In addition, we introduce a library of MIP instances with embedded ML constraints. The project is available at <https://github.com/Opt-Mucca/PySCIPOpt-ML>.

## 1 Introduction and Related Work

Many problems coming from real-world applications can be modelled using *Mixed-Integer Programming (MIP)*. They can be expressed as

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & f_0(x) \\ \text{s.t.} \quad & f_i(x) \leq 0, \forall i \in [m], \\ & x_i \in \mathbb{Z}, \forall i \in \mathcal{I}, \end{aligned} \tag{P}$$

where each  $f_i : \mathbb{R}^n \rightarrow \mathbb{R}, i = 0, \dots, m$  is continuous and  $\mathcal{I} \subseteq [n]$  denotes the index set of the integer-constrained variables. Due to its ability to represent complex systems and decision-making processes across various industries, MIP has become a standard tool to solve optimisation problems arising in real-world applications.

A common challenge of finding a suitable MIP formulation (P) of a system is the presence of unknown or highly complex relationships. This poses a problem since traditional MIP approaches rely on precise definitions of constraints and the objective function. To address this challenge, *Machine Learning (ML)* predictors are often used to approximate these relationships and are then embedded in (P) as surrogate models. To embed a trained ML predictor  $g : \mathbb{R}^n \rightarrow \mathbb{R}$ , it must first be formulated as a MIP, before it can be added to (P). However, with the growing popularity and accessibility of different frameworks [1, 26, 10, 16, 25], automatically formulating trained ML models coming from these different architectures into MIPs has become a highly non-trivial, yet necessary task.

**Contribution.** In this paper, we present PySCIPOpt-ML, an open-source tool for the automatic formulation of trained ML predictors as MIPs. By directly interfacing with a broad range of commonly used ML

---

\*Zuse Institute Berlin, Department of Mathematical Optimization, Takustr. 7, 14195 Berlin

<sup>†</sup>Chair of Software and Algorithms for Discrete Optimization, Institute of Mathematics, Technische Universität Berlin, Straße des 17. Juni 135, 10623 Berlin, Germany

<sup>‡</sup>Gurobi GmbH, Ulmenstr. 37-39, 60325 Frankfurt am Main, Germany




Figure 1: How trained ML predictors can be automatically formulated as MIPs and embedded in an optimisation problem using our Python package PySCIPOpt-ML.

frameworks, PySCIPOpt-ML allows for the easy integration of ML constraints into MIPs, which can then be solved by the state-of-the-art open-source solver SCIP [6]. The package supports various models from Scikit-Learn [26], XGBoost [10], LightGBM [16], and PyTorch [25], making it a general tool to easily optimise MIPs with embedded ML constraints (see Figure 1). To represent predictors as MIPs, we use formulations that prioritise numerical stability to enhance solution robustness. In addition, we introduce an expandable library of MIP instances with embedded ML constraints based on real-world data.

To summarise, our contributions are the following:

1. We introduce the **Python package PySCIPOpt-ML**<sup>4</sup>, which **directly interfaces various ML frameworks with the open-source solver SCIP** (see Section 2).
2. We utilise **MIP formulations prioritising numerical stability**, especially when solving with SCIP (see Section 3).
3. We introduce a **new library of MIP instances with embedded ML constraints called SurrogateLIB**<sup>5</sup>, as well as the corresponding instance generators (see Section 4).

**Related Work.** In recent years, more and more industries have leveraged ML frameworks to approximate unknown or complex relationships in optimisation problems. For instance, ML predictors have been used to optimise cancer treatment plans [5], minimise energy consumption [22], and in the design of energy systems [15]. Embedded ML predictors have also been used in the optimisation of gas production and water management systems [20], as well as in the optimisation of retail prices [11] and locations [14], to name only a few examples. For surveys of embedded ML models through surrogate optimisation, see [18, 7, 21].

The increasing demand for embedding ML predictors into MIPs has motivated the development of various tools for automatic integration. The commercial MIP solver Gurobi released the Python package Gurobi Machine Learning<sup>6</sup> [24], which integrates trained regression models into optimisation problems that can then be solved with Gurobi. Another package that relies on Gurobi is Janos [4], which supports linear and logistic regression, as well as neural networks with rectified linear unit (ReLU) activation functions. ENTMOOT [28] focuses on representing gradient-boosted tree models from LightGBM. On the other hand, reluMIP [19] allows the user to integrate ReLU neural networks trained with TensorFlow [1] within optimisation problems modelled with either Pyomo [8] or Gurobi. OMLT<sup>7</sup> [9] focuses on embedding neural networks and gradient-boosted trees via the general MIP modelling framework Pyomo [8] and the general purpose ML interface ONNX<sup>8</sup>.

While many of these packages focus on supporting specific types of predictors, PySCIPOpt-ML allows the user to embed a wide variety of predictors. Specifically, for both regression and classification tasks, neural networks with various activation functions, linear and logistic regression models, decision trees, gradient boosted trees, and random forests.

<sup>4</sup><https://github.com/Opt-Mucca/PySCIPOpt-ML>

<sup>5</sup><https://zenodo.org/records/10357875>

<sup>6</sup><https://github.com/Gurobi/gurobi-machinelearning>

<sup>7</sup><https://github.com/cog-imperial/OMLT>

<sup>8</sup><https://github.com/onnx/onnx>

ML Predictor	PyTorch	Scikit-Learn	LightGBM	XGBoost
Decision Trees	-	R+C	R+C	R+C
Gradient Boosted Trees	-	R+C	R+C	R+C
Random Forests	-	R+C	R+C	R+C
Neural Networks	R+C	R+C	-	-
Linear Regression	-	R	-	-
Logistic Regression	-	R+C	-	-
Partial Least Squares	-	R	-	-

Table 1: Overview of ML predictors coming from different ML frameworks supported by PySCIPOpt-ML. “R” indicates that (multi-)regression is supported and “C” denotes the support of (multi-)classification. “-” show predictors not supported by a framework and thus not supported by PySCIPOpt-ML.

## 2 About the Package

PySCIPOpt-ML directly interfaces with the popular ML frameworks Scikit-Learn [26], XGBoost [10], LightGBM [16], and PyTorch [25]. Each framework allows the user to train ML predictors of different types, all of which can be embedded into a MIP by our package. A detailed overview of what PySCIPOpt-ML supports is given in Table 1.

An advantage of PySCIPOpt-ML is that it seamlessly joins all parts of the modelling, training, and optimisation processes. At no point is there a need to export the ML predictor or MIP model into some intermediary language. A central function (`add_predictor_constr()`) acts as the main interface between the ML frameworks training the predictor and the solver SCIP solving the problem. To create and solve a MIP problem with an embedded ML predictor, the user must only do the following:

1. Define the optimisation problem, `scip_model`, by first adding the explicit constraints and objective function.
2. Given data (`input_data, output_data`), choose which predictor should be trained and embed the trained predictor automatically in the optimisation problem as a constraint. This is done by simply calling:

```
add_predictor_constr(scip_model, predictor, input_data, output_data).
```

3. Solve the updated `scip_model` with SCIP.

In summary, PySCIPOpt-ML helps to easily create, embed, and solve optimisation problems with ML surrogate models and spares the user from navigating multiple platforms at the same time.

## 3 MIP Formulations of ML Predictors

To embed an ML predictor in a MIP, we first need to represent the predictor itself as a MIP. The formulation of the predictor has a significant effect on the ability to efficiently optimise the larger optimisation problem. There has been a substantial amount of work on determining the best performing formulations for popular ML predictors. In particular, neural networks with ReLU activation functions [13, 27, 3, 17] and ensemble based methods that use decision trees [23, 28, 2, 24], have received a lot of attention in the past. Often, the most natural formulation of these predictors involves quadratic constraints. Due to performance issues, however, formulations that utilise big-M, indicator, and SOS constraints are in general preferred.

### 3.1 Neural Networks

The current standard MIP formulation for feed forward neural networks with ReLU activation functions is to use big-M constraints [13, 3, 9]. These formulations are prone to numerical issues when the bounds on the input are large, and thus benefit greatly from tighter bounds. For this reason, either feasibility based or optimality based bound tightening is typically performed prior to optimisation [13, 9]. It is therefore important that the user specifies bounds on the variables of the wider optimisation problem. Since this is not always possible, PySCIPOpt-ML utilises a SOS1-based formulation to represent ReLU activation functions. By doing so, we aim to maintain numerical stability even in the absence of tight variable bounds, as the validity of the constraints can be enforced by branching [12].

Consider a single neuron in a fully connected layer  $k$ . Let  $x \in \mathbb{R}^n$  be the output of the previous layer  $k - 1$ ,  $y \in \mathbb{R}$  be the output of the given neuron,  $w \in \mathbb{R}^n$  and  $b \in \mathbb{R}$  be the weights and bias connecting the previous layer to the neuron, and  $s \in \mathbb{R}_{\geq 0}$  be some slack variable. We use the following formulation of the ReLU activation at the neuron:

$$y = \sum_{i=1}^n w_i x_i + b + s,$$

$$\text{SOS1}(s, y).$$

An SOS1 constraint enforces that at most one variable in the constraint can be non-zero. The above formulation ensures that the slack is only non-zero when  $\sum_{i=1}^n w_i x_i + b$  is negative. In this case, the slack must also take exactly the value of  $-b - \sum_{i=1}^n w_i x_i$ , as the SOS1 constraint then forces  $y$  to 0.

As PySCIPOpt-ML uses SCIP [6] as a solver, other activation functions such as sigmoid and tanh can be added directly.

### 3.2 Tree-Based Models

To model decision trees, PySCIPOpt-ML uses the formulation first introduced in Gurobi Machine Learning [24]. As opposed to quadratic formulations [23] and linear big-M formulations [2], the Gurobi Machine Learning formulation uses a series of indicator constraints for each leaf node ensuring that any input maps to the appropriate leaf node of the tree. A more detailed overview of this formulation is available at<sup>9</sup>. The MIP formulations of gradient boosted trees and random forests are simply a linear combination of the individual decision tree formulations.

### 3.3 Argmax Formulation

For many classification tasks, the ML predictor’s evaluation is determined by an argmax function call. That is, if each class  $j$  from the index set  $\mathcal{J}$  has some predicted score  $y_j \in \mathbb{R}$ , the predictor outputs the class with the highest score. To model this functionality with MIP, we define a variable  $m \in \mathbb{R}$ , binary variables  $z_j \in \{0, 1\}$ ,  $j \in \mathcal{J}$ , and slack variables  $s_j \geq 0$ ,  $j \in \mathcal{J}$ . The formulation is:

$$y_j + s_j - m = 0, \quad \forall j \in \mathcal{J}$$

$$\text{SOS1}(z_j, s_j), \quad \forall j \in \mathcal{J}$$

$$\sum_{i=1}^{|\mathcal{J}|} z_j = 1$$

In the above formulation, the variable  $m$  implicitly represents  $\max\{y_j : j \in \mathcal{J}\}$ : If  $m$  is less than the maximum, there exists a  $j' \in \mathcal{J}$  s.t.  $y_{j'} \geq m$ , and for which  $y_{j'} + s_{j'} - m = 0$  cannot be valid. On the other hand, if  $m$  is greater than the maximum, then all slack variables are non-zero, and  $\sum_{i=1}^{|\mathcal{J}|} z_j = 1$  cannot be satisfied due to the SOS1 constraints. Because  $m$  is the maximum, exactly one  $z_{j'}$  is set to 1, where  $y_{j'} = \max\{y_j : j \in \mathcal{J}\}$ .

To truly return the argmax instead of a binary variable corresponding to the location of the maximum value, one can introduce a variable  $a \in \mathbb{Z}$ , and the constraint:

$$a = \sum_{i=1}^{|\mathcal{J}|} j * z_j.$$

The variable  $a$  will then have value corresponding to that returned by the traditional argmax function.

## 4 Instance Library

Accompanying PySCIPOpt-ML, we present a new instance library SurrogateLib<sup>10</sup> consisting of MIPs that contain embedded ML predictors as either constraints or components of the objective function. These instances are inspired by diverse real-world applications, including auto manufacturing (see Section 4.1),

<sup>9</sup><https://gurobi-machinelearning.readthedocs.io/en/stable/user/mip-models.html>

<sup>10</sup><https://zenodo.org/records/10357875>

water treatment, and adversarial attacks. By avoiding synthetically generated data to train the ML predictors, we aim to showcase the practical benefits of embedding predictive models in optimisation problems.

Alongside the library, we also include instance generators of the introduced problems that allow users to create homogeneous instances of varying difficulty. Thus, PySCIPOpt-ML and SurrogateLib present an opportunity to generate semi-realistic MIPs of controllable complexity by adjusting the size of the embedded ML predictors as well as the parameters of the instances. This way, we provide the MIP community with a set of homogeneous instances that are sufficiently diverse and relatively easy to solve without being trivial. In the following, we will present one of SurrogateLib’s base problems.

#### 4.1 Auto Manufacturer Example

Let us take the point of view of an auto manufacturer that needs to design a new vehicle. Based on historical data about past sales, the manufacturer aims to maximise the number of vehicles sold, while ensuring that the designed vehicle is sufficiently different to any other popular vehicle already on the market. Additionally, the manufacturer wants to ensure that the designed vehicle has a high resell value and is fuel-efficient.

Let  $x \in \mathbb{R}^n$  be the vector of variable vehicle features that the manufacturer can control. In our case, we have  $n = 10$  features: vehicle type, engine size, horsepower, wheelbase, width, length, curb weight, fuel capacity, efficiency and power performance factor. Every feature has a lower bound  $l_i$  and an upper bound  $u_i$ , that is,  $l \leq x \leq u$  with  $l, u \in \mathbb{R}^n$ .

Furthermore, let  $z_{\text{price}}$ ,  $z_{\text{resell}}$ , and  $z_{\text{sold}}$  be the features of the vehicle that the manufacturer cannot control. In this example, these uncontrollable features are the price of the vehicle, its resell value, and the number of sold vehicles.

Since the relationship between  $x$  and  $(z_{\text{price}}, z_{\text{resell}}, z_{\text{sold}})$  is unknown, the auto manufacturer uses ML to learn from real-world historic data<sup>11</sup>. To do so, a ML predictor,  $g$ , is trained that takes as input the controllable features of the vehicle and outputs the predicted price. Two more ML predictors,  $f$  and  $h$ , are trained, which take as input the controllable features and the predicted price of the vehicle, and output the predicted amount of vehicles sold and predicted resell value respectively.

Recall that the manufacturer aims to design a vehicle different to what already exists, requiring us to introduce additional constraints. For a finite index set  $\mathcal{V}$ , we denote by  $\bar{x}^j \in \mathbb{R}^{10}$ ,  $j \in \mathcal{V}$ , the controllable features of the  $|\mathcal{V}|$  many most popular vehicles. To ensure that  $x$  is sufficiently different to each  $\bar{x}^j$ , we consider the relative difference of each feature and enforce that the cumulative difference among all features are above a given threshold  $\gamma_j$ . Formally, we obtain the constraints

$$\sum_{i=1}^{10} \frac{|x_i - \bar{x}_i^j|}{u_i - l_i} \geq \gamma_j, \quad \forall j \in \mathcal{V}. \quad (1)$$

Note that even though the absolute value function is non-linear, it can be easily modelled using linear constraints and binary variables.

Finally, let  $\alpha$  be the minimum ratio of the resell value compared to the actual purchase price, and let  $\beta$  be the minimum level of fuel-efficiency. To ensure that the resulting vehicle has a high enough resell value and is fuel-efficient enough, we need to add the constraints:

$$z_{\text{resell}} \geq \alpha z_{\text{price}}, \quad (2)$$

$$x_{\text{eff}} \geq \beta. \quad (3)$$

Here,  $x_{\text{eff}}$  denotes the entry in  $x$  that corresponds to the fuel efficiency feature.

The resulting optimisation problem with embedded ML predictors is then given by:

$$\begin{aligned} \max_x \quad & z_{\text{sold}} \\ \text{s.t.} \quad & z_{\text{price}} = g(x), \\ & z_{\text{sold}} = f(x, z_{\text{price}}), \\ & z_{\text{resell}} = h(x, z_{\text{price}}), \\ & (1) - (3), \\ & l \leq x \leq u. \end{aligned}$$

<sup>11</sup><https://www.kaggle.com/datasets/gagandeep16/car-sales>

This parameterised problem can now be used to generate a set of homogeneous instances. By choosing different parameters  $\alpha, \beta, \gamma_j$  and  $\bar{x}_j$  for  $j \in \mathcal{V}$  as well as different ML models to obtain predictors  $f, g$  and  $h$ , we are able to create instances of varying complexity while still having the same underlying structure.

## 5 Conclusion

In this paper we have introduced the Python package PySCIPOpt-ML, which helps to automatically embed trained ML predictors in MIPs. We have briefly introduced the API, expanded on some of the core MIP formulations, and walked through a real-world example. In addition, we have introduced a new set of homogeneous semi-realistic MIP instances that form the extendable library SurrogateLib.

## Acknowledgements

We thank Mohammed Ghannam, Robert Luce, and Franziska Schlösser for their help with deploying SCIP via Pip. The work for this article has been conducted in the Research Campus MODAL funded by the German Federal Ministry of Education and Research (BMBF) (fund numbers 05M14ZAM, 05M20ZBM).

## References

- [1] Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems (2015), <https://www.tensorflow.org/>, software available from tensorflow.org
- [2] Ammari, B.L., Johnson, E.S., Stinchfield, G., Kim, T., Bynum, M., Hart, W.E., Pulsipher, J., Laird, C.D.: Linear Model Decision Trees as Surrogates in Optimization of Engineering Applications. *Computers & Chemical Engineering* **178** (2023). <https://doi.org/https://doi.org/10.1016/j.compchemeng.2023.108347>
- [3] Anderson, R., Huchette, J., Ma, W., Tjandraatmadja, C., Vielma, J.P.: Strong mixed-integer programming formulations for trained neural networks. *Mathematical Programming* **183**(1-2), 3–39 (2020)
- [4] Bergman, D., Huang, T., Brooks, P., Lodi, A., Raghunathan, A.U.: JANOS: an integrated predictive and prescriptive modeling framework. *INFORMS Journal on Computing* **34**(2), 807–816 (2022)
- [5] Bertsimas, D., O’Hair, A., Relyea, S., Silberholz, J.: An Analytics Approach to Designing Combination Chemotherapy Regimens for Cancer. *Management Science* **62**(5), 1511–1531 (2016)
- [6] Bestuzheva, K., Besançon, M., Chen, W.K., Chmiela, A., Donkiewicz, T., van Doornmalen, J., Eifler, L., Gaul, O., Gamrath, G., Gleixner, A., Gottwald, L., Graczyk, C., Halbig, K., Hoen, A., Hojny, C., van der Hulst, R., Koch, T., Lübbecke, M., Maher, S.J., Matter, F., Mühmer, E., Müller, B., Pfetsch, M.E., Rehfeldt, D., Schlein, S., Schlösser, F., Serrano, F., Shinano, Y., Sofranac, B., Turner, M., Vigerske, S., Wegscheider, F., Wellner, P., Weninger, D., Witzig, J.: Enabling Research through the SCIP Optimization Suite 8.0. *ACM Trans. Math. Softw.* **49**(2) (2023)
- [7] Bhosekar, A., Ierapetritou, M.: Advances in surrogate based modeling, feasibility analysis, and optimization: A review. *Computers & Chemical Engineering* **108**, 250–267 (2018)
- [8] Bynum, M.L., Hackebeil, G.A., Hart, W.E., Laird, C.D., Nicholson, B.L., Siirola, J.D., Watson, J.P., Woodruff, D.L.: *Pyomo—optimization modeling in Python*, vol. 67. Springer Science & Business Media, third edn. (2021)
- [9] Ceccon, F., Jalving, J., Haddad, J., Thebelt, A., Tsay, C., Laird, C.D., Misener, R.: OMLT: Optimization & machine learning toolkit. *The Journal of Machine Learning Research* **23**(1), 15829–15836 (2022)
- [10] Chen, T., Guestrin, C.: XGBoost: A Scalable Tree Boosting System. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. pp. 785–794. KDD ’16, ACM (2016)
- [11] Ferreira, K., Lee, B., Simchi-Levi, D.: Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management* **18** (1) (2015)



- [12] Fischer, T., Pfetsch, M.E.: Branch-and-cut for linear programs with overlapping sos1 constraints. *Mathematical Programming Computation* **10**, 33–68 (2018)
- [13] Grimstad, B., Andersson, H.: ReLU networks as surrogate models in mixed-integer linear programs. *Computers & Chemical Engineering* **131**, 106580 (2019)
- [14] Huang, T., Bergman, D., Gopal, R.: Predictive and prescriptive analytics for location selection of add-on retail products. *Production and Operations Management* **28** (4) (2018)
- [15] Jalving, J., Ghouse, J., Cortes, N., Gao, X., Knueven, B., Agi, D., Martin, S., Chen, X., Guittet, D., Tumbalam-Gooty, R., Bianchi, L., Beattie, K., Gunter, D., Siirola, J.D., Miller, D.C., Dowling, A.W.: Beyond price taker: Conceptual design and optimization of integrated energy systems using machine learning market surrogates. *Applied Energy* **351**, 121767 (2023)
- [16] Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., Liu, T.Y.: LightGBM: A highly efficient gradient boosting decision tree. *Advances in Neural Information Processing Systems* **30**, 3146–3154 (2017)
- [17] Kronqvist, J., Misener, R., Tsay, C.: Between steps: Intermediate relaxations between big-M and convex hull formulations. In: *International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research*. pp. 299–314. Springer (2021)
- [18] Lombardi, M., Milano, M.: Boosting combinatorial problem modeling with machine learning. *arXiv preprint arXiv:1807.05517* (2018)
- [19] Lueg, L., Grimstad, B., Mitsos, A., Schweidtmann, A.M.: reluMIP: Open Source Tool for MILP Optimization of ReLU Neural Networks (2021). <https://doi.org/https://doi.org/10.5281/zenodo.5601907>, [https://github.com/ChemEngAI/ReLU\\_ANN\\_MILP](https://github.com/ChemEngAI/ReLU_ANN_MILP)
- [20] López-Flores, F., Lira-Barragn, L., Rubio-Castro, E., El-Halwagi, M., Ponce-Ortega, J.: Hybrid machine learning-mathematical programming approach for optimizing gas production and water management in shale gas fields. *ACS Sustainable Chemistry & Engineering* **11** (2023)
- [21] McBride, K., Sundmacher, K.: Overview of surrogate modeling in chemical process engineering. *Chemie Ingenieur Technik* **91**(3), 228–239 (2019)
- [22] Misaghian, M.S., Tardioli, G., Cabrera, A.G., Salerno, I., Flynn, D., Kerrigan, R.: Assessment of carbon-aware flexibility measures from data centres using machine learning. *IEEE Transactions on Industry Applications* **59**(1), 70–80 (2023)
- [23] Mistry, M., Letsios, D., Krennrich, G., Lee, R.M., Misener, R.: Mixed-integer convex nonlinear optimization with gradient-boosted trees embedded. *INFORMS Journal on Computing* **33**(3), 1103–1119 (2021)
- [24] Optimization, G.: Gurobi-machinelearning (2023), <https://github.com/Gurobi/gurobi-machinelearning/>
- [25] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: PyTorch: An Imperative Style, High-Performance Deep Learning Library. *Advances in Neural Information Processing Systems* **32**, 8024–8035 (2019)
- [26] 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.: Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011)
- [27] Schweidtmann, A.M., Mitsos, A.: Deterministic global optimization with artificial neural networks embedded. *Journal of Optimization Theory and Applications* **180**(3), 925–948 (2019)
- [28] Thebelt, A., Kronqvist, J., Mistry, M., Lee, R.M., Sudermann-Merx, N., Misener, R.: ENTMOOT: a framework for optimization over ensemble tree models. *Computers & Chemical Engineering* **151**, 107343 (2021)