

The Machine Learning for Combinatorial Optimization Competition (ML4CO): Results and Insights

Co-organizers

Maxime Gasse*, Simon Bowly, Quentin Cappart, Jonas Charfreitag,
Laurent Charlin, Didier Chételat, Antonia Chmiela, Justin Dumouchelle,
Ambros Gleixner, Aleksandr M. Kazachkov, Elias Khalil, Pawel Lichocki,
Andrea Lodi, Miles Lubin, Chris J. Maddison, Christopher Morris,
Dimitri J. Papageorgiou, Augustin Parjadis, Sebastian Pokutta,
Antoine Prouvost, Lara Scavuzzo, Giulia Zarpellon

Primal task winners

Linxin Yang	YANGLINXIN@CUHK.EDU.CN
Sha Lai	221049039@LINK.CUHK.EDU.CN
Akang Wang	WANGAKANG@SRIBD.CN
Xiaodong Luo	XIAODONGLUO@CUHK.EDU.CN
Xiang Zhou	ZHOUXIANG60@HUAWEI.COM
Haohan Huang	HUANGHAOHAN@HUAWEI.COM
Shengcheng Shao	SHAOSHENGCHENG@HUAWEI.COM
Yuanming Zhu	ZHUYUANMING5@HUAWEI.COM
Dong Zhang	ZHANGDONG48@HUAWEI.COM
Tao Quan	QUANTAO@HUAWEI.COM

Dual task winners

Zixuan Cao	CAOZIXUAN.PERCY@STU.PKU.EDU.CN
Yang Xu	1800010740@PKU.EDU.CN
Zhewei Huang	HUANGZHEWEI@MEGVII.COM
Shuchang Zhou	ZSC@MEGVII.COM

Configuration task winners

Chen Binbin	CBB18@MAILS.TSINGHUA.EDU.CN
He Minggui	HEMINGGUI@HUAWEI.COM
Hao Hao	52194506007@STU.ECNU.EDU.CN
Zhang Zhiyu	ZHANGZHIYU6@HUAWEI.COM
An Zhiwu	ANZHIWU1@HUAWEI.COM
Mao Kun	MAOKUN@HUAWEI.COM

Editor: Douwe Kiela, Marco Ciccone, Barbara Caputo

* Primary contact (maxime.gasse@polymtl.ca)

Abstract

Combinatorial optimization is a well-established area in operations research and computer science. Until recently, its methods have focused on solving problem instances in isolation, ignoring that they often stem from related data distributions in practice. However, recent years have seen a surge of interest in using machine learning as a new approach for solving combinatorial problems, either directly as solvers or by enhancing exact solvers. Based on this context, the ML4CO aims at improving state-of-the-art combinatorial optimization solvers by replacing key heuristic components. The competition featured three challenging tasks: finding the best feasible solution, producing the tightest optimality certificate, and giving an appropriate solver configuration. Three realistic datasets were considered: *balanced item placement*, *workload apportionment*, and *maritime inventory routing*. This last dataset was kept anonymous for the contestants.

Keywords: Combinatorial optimization, machine learning for combinatorial optimization

1. Introduction

The *Machine Learning for Combinatorial Optimization* competition (ML4CO) aims at improving state-of-the-art combinatorial optimization solvers by replacing key heuristic components with machine learning models. The main scientific question is the following: “*Is machine learning a viable option for improving traditional combinatorial optimization solvers on specific problem distributions, when historical data is available?*”

While most combinatorial optimization solvers are presented as general-purpose, *one-size-fits-all* algorithms, this competition focuses on the design of application-specific algorithms from historical data. This general problem captures a highly practical scenario relevant to many application areas, where a practitioner repeatedly solves problem instances from a specific distribution with similar patterns and characteristics. For example, managing a large-scale energy distribution network requires solving similar CO problems daily, with a fixed power grid structure while only the demand changes over time. This demand change is hard to capture by hand-engineered expert rules, and ML-enhanced approaches offer a possible solution to detect typical patterns in the demand history. Other examples include crew scheduling problems that have to be solved daily or weekly with minor variations or vehicle routing where the traffic conditions change over time, but the overall transportation network does not.

The competition features three challenges for machine learning. Each of them corresponds to a specific control task arising in the open-source solver *SCIP* (Gamrath et al., 2020) and is exposed through a unified OpenAI Gym-like API based on the Python library *Ecole* (Prouvost et al., 2020). The three challenges are as follows: (1) *a primal task*, consisting of producing the best feasible solution, (2) *a dual task*, consisting of producing the best branching decisions, and (3) *a configuration task*, consisting of finding the best parameters before calling the solver. For each challenge, participants were evaluated on three problem benchmarks originating from diverse application areas, each represented as a collection of *mixed-integer linear program* (MILP) instances.

2. Datasets

The participants’ solutions were evaluated on three problem benchmarks from diverse application areas for each challenge. Each participant submitted a decision-making code, i.e., an algorithmic solution or trained ML model, for each of the benchmarks, or a single code that works for all

benchmarks. A problem benchmark consists of a collection of MILP instances in the standard MPS file format. Each benchmark was split into a training and a test set, originating from the same problem distribution. While the training instances were made public at the beginning of the competition for participants to train their models, the test instances were kept hidden for evaluation purposes and were only revealed at the end of the competition. The first two problem benchmarks were inspired by real-life applications of large-scale systems at Google, while the third benchmark was presented to the participants as an anonymous problem inspired by a real-world, large-scale industrial application. The dataset used are publicly available on Github.¹

Benchmark 1: Balanced Item Placement This problem involves spreading items, e.g., files or processes, across containers, e.g., disks or machines, utilizing them evenly. Items can have multiple copies, but at most, one copy can be placed in a single bin. The number of items that can be moved is constrained, modeling the real-life situation of a live system for which some placement already exists. Each problem instance is modeled as a MILP, using a multi-dimensional, multi-knapsack formulation. This dataset contains 10000 training instances (pre-split into 9900 train and 100 validation instances).

Benchmark 2: Workload Apportionment This problem deals with apportioning workloads, e.g., data streams, across as few workers, e.g., servers, as possible. The apportionment is required to be robust to any worker’s failure. Each instance problem is modeled as a MILP, using a bin-packing with an apportionment formulation. This dataset contains 10000 training instances (pre-split into 9900 train and 100 validation instances).

Benchmark 3: Maritime Inventory Routing (Anonymous Problem) This problem plays an integral role in global bulk shipping. The instances corresponding to this benchmark are assembled from a public dataset (Papageorgiou et al., 2014), whose origin was kept secret to prevent cheating. Reverse-engineering for the purpose of recovering the test set was forbidden. The dataset contains 118 training instances (pre-split into 98 train and 20 validation instances).

3. Evaluation Metrics

Each of the three challenges is associated with a specific evaluation metric reflecting a different objective. We describe how each metric is computed over a single problem instance. The final goal of the participants is to optimize this metric in expectation over a hidden collection of test instances. Because the evaluation metrics are time-dependent, all evaluations were run on the same hardware setup, using an Intel Xeon processor with 2.4GHz, 20GB of RAM, and an Nvidia Tesla V100-8G GPU with 8GB of GPU memory. A maximum time budget was given for each task to process each test instance (5, 15, and 15 minutes for the primal, dual, and configuration tasks, respectively), after which the environment was terminated. By doing so, participants were asked to focus on making both good and fast decisions.

In general, a MILP instance is expressed as follows,

$$\begin{aligned} & \arg \min_{\mathbf{x}} \quad \mathbf{c}^\top \mathbf{x} \\ \text{subject to} \quad & \mathbf{A}^\top \mathbf{x} \leq \mathbf{b} \\ & \mathbf{x} \in \mathbb{Z}^p \times \mathbb{R}^{n-p}, \end{aligned}$$

1. <https://github.com/ds4dm/ml4co-competition>

where $\mathbf{c} \in \mathbb{R}^n$ denotes the coefficients of the linear objective, $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$, denote the coefficients and upper bounds of the linear constraints, respectively, while n is the total number of variables, $p \leq n$ is the number of integer-constrained variables, and m is the number of linear constraints. We used the following three metrics for the evaluation.

Primal Integral This metric measures the area under the curve of the solver’s primal bound, i.e., its global upper bound, which corresponds to the value of the best feasible solution found so far. By providing better feasible solutions over time, the value of the primal bound decreases. With a time limit T , the primal integral is expressed as follows,

$$\int_{t=0}^T \mathbf{c}^\top \mathbf{x}_t^* dt - T \mathbf{c}^\top \mathbf{x}^*,$$

where \mathbf{x}_t^* is the best feasible solution found at time t , so that $\mathbf{c}^\top \mathbf{x}_t^*$ is the primal bound at time t , and $T \mathbf{c}^\top \mathbf{x}^*$ is an instance-specific constant that depends on the optimal solution \mathbf{x}^* . The primal integral is to be *minimized*, and takes an optimal value of 0. To compute this metric unambiguously, a trivial initial solution x_0^* is always provided to the solver at the beginning of the solving process. Also, the constant term $\mathbf{c}^\top \mathbf{x}^*$ can be safely ignored at training time when participants train their control policy. However, when we evaluated the participant submissions at test time, this constant term, or a proper substitute, was incorporated in the reported metric.

Dual Integral This metric measures the area over the curve of the solver’s dual bound, i.e., its global lower bound, which usually corresponds to a solution of a valid relaxation of the MILP. By branching, the linear relaxation corresponding to the branch-and-bound tree’s leaves gets tightened, and the dual bound increases over time. With a time limit T , the dual integral is expressed as follows,

$$T \mathbf{c}^\top \mathbf{x}^* - \int_{t=0}^T \mathbf{z}_t^* dt,$$

where \mathbf{z}_t^* is the best dual bound at time t , and $T \mathbf{c}^\top \mathbf{x}^*$ is an instance-specific constant that depends on the optimal solution value $\mathbf{c}^\top \mathbf{x}^*$. The dual integral is to be *minimized*, and takes an optimal value of 0.

In the context of branching, this metric is unambiguous to compute, as the root node of the tree always provides an initial dual bound \mathbf{z}_0^* at the beginning of branching. The constant term $\mathbf{c}^\top \mathbf{x}^*$ can be safely ignored for training, but it was incorporated in the evaluation metric.

Primal-Dual Gap Integral This metric measures the area between two curves, the solver’s *primal bound* and *dual bound*. Hence, the metric benefits both from improvements obtained from the primal side (finding good feasible solutions), and on the dual side (producing a tight optimality certificate). With a time limit T , the primal-dual gap integral is expressed as follows,

$$\int_{t=0}^T \mathbf{c}^\top \mathbf{x}_t^* - \mathbf{z}_t^* dt.$$

The primal-dual gap integral is to be *minimized*, and takes an optimal value of 0. In the context of algorithm configuration, an initial value is required for the two curves at time $t = 0$. Therefore, an initial trivial solution x_0^* and a valid initial dual bound \mathbf{z}_0^* are always provided to the solver for this task.

4. Solving the Primal Task (Winning Solution by CUHKSZ_ATD)

The primal task deals with finding good primal solutions at the root node of the branch-and-bound tree (Khalil et al., 2017; Nazari et al., 2018; Nair et al., 2020). To that end, the environment (SCIP solver) does not perform any branching but enters an infinite loop at the root node, which collects the solutions proposed by the agents, evaluates their feasibility and updates the overall best solution reached so far, thus lowering the current primal bound (upper bound). The metric of interest for this task is the primal integral, which considers the rate at which the primal bound decreases over time. To model a realistic scenario, each problem instance has been preprocessed by SCIP (problem reduction, cutting planes, etc.). Moreover, the root linear program (LP) relaxation was solved before the participants were asked to produce feasible solutions. To prevent SCIP from searching for primal solutions by itself, all primal heuristics were deactivated. Further, to compute this metric unambiguously, even when no solution has been found yet, an initial primal bound (trivial solution value) was provided for each instance, which is to be given to SCIP at the beginning of the solving process in the form of a user objective limit. Execution time was limited to five minutes. We exploit the problem structures and tackle item placement, workload apportionment, and anonymous problems by utilizing classic primal heuristics in a more judicious manner.

4.1. Balanced Item Placement

Let I denote the set of items and J denote the set of containers. Let a binary variable x_{ij} be 1 if item i is placed in container j and 0 otherwise. Each item will be placed in exactly a single container, as shown by constraints (2). Let K represent the set of dimensions. For dimension $k \in K$ of container $j \in J$, knapsack constraints (3) represent some physical considerations while (4) and (5) properly account for the placement unevenness, which is penalized in the objective (1).

$$\min_{x,y,z} \quad \sum_{j \in J} \sum_{k \in K} \alpha_k y_{jk} + \sum_{k \in K} \beta_k z_k \quad (1)$$

$$\text{s.t.} \quad \sum_{j \in J} x_{ij} = 1 \quad \forall i \in I \quad (2)$$

$$\sum_{i \in I} a_{ik} x_{ij} \leq b_k \quad \forall j \in J, \forall k \in K \quad (3)$$

$$\sum_{i \in I} d_{ik} x_{ij} + y_{jk} \geq 1 \quad \forall j \in J, \forall k \in K \quad (4)$$

$$y_{jk} \leq z_k \quad \forall j \in J, \forall k \in K \quad (5)$$

$$x_{ij} \in \{0, 1\} \quad \forall i \in I, \forall j \in J \quad (6)$$

$$y_{jk} \geq 0 \quad \forall j \in J, \forall k \in K \quad (7)$$

We analyze 10,000 item placement instances and find out that: (i) $|I| = 105, |J| = 10$; (ii) a_{ik}, d_{ik} values of 5 items are big and placing any two of them in the same container would incur large penalty. Based on this empirical finding, we place the 5 big items into the first five containers respectively before applying any primal heuristics.

Meta-heuristics We apply a greedy method in which items are first sorted based on their sizes and then assigned to containers. This will produce the very first feasible solution. After that, we select one or two items respectively from two containers and swap them if this leads to a better incumbent.

Math-heuristics We consider a construction method and an improvement method, based on solving mathematical models. The construction method consists of two steps: (i) first assign items to the first five containers by solving an assignment model; (ii) then assign the remaining items to the last five containers by solving a sub-MIP. For solution improvement, we properly choose two out of the last five containers, and then solve a sub-MIP to reassign items within those two containers optimally.

4.2. Workload Apportionment

Let M and N denote the set of tasks and the set of machines, respectively. For task $i \in M$, only a subset of machines, denoted by $N^i \subseteq N$, are accessible. Let a binary variable y_j be 1 if machine j is used and 0 otherwise. Let x_{ij} denote the amount of workload from task i to machine j , as defined in constraints (9). Constraints (10) enforce the capacity requirement for each machine. The apportionment is required to be robust to any one machine’s failure, as indicated by constraints (11).

$$\min_{x,y} \quad \sum_{j \in N} y_j \quad (8)$$

$$\text{s.t.} \quad x_{ij} \leq a_i y_j \quad \forall i \in M, \forall j \in N^i \quad (9)$$

$$\sum_{i \in M: j \in N^i} x_{ij} \leq b_j \quad \forall j \in N \quad (10)$$

$$\sum_{j \in N^i \setminus \{j'\}} x_{ij} \geq a_i \quad \forall i \in M, \forall j' \in N^i \quad (11)$$

$$y_j \in \{0, 1\} \quad \forall j \in N \quad (12)$$

$$0 \leq x_{ij} \leq b_j \quad \forall i \in M, \forall j \in N^i \quad (13)$$

Rounding up a solution to the linear programming relaxation of model (8) – (13) would produce a feasible solution to (8) – (13). To further exploit the possibility of rounding a fractional solution towards a new incumbent, we choose a rounding threshold parameter η in an adaptive manner and round up y_j only if it exceeds η . Specifically, we first select a target objective value based on the current primal and dual bounds and then determine η via quantile selection such that after rounding with η the objective matches the pre-determined value. If the rounding step produces a new incumbent solution, we then update the primal bound; otherwise, we set the dual bound to the corresponding objective value. We iterate this process until the primal-dual gap falls below a predetermined value. In model (8) – (13), constraints (10) can be tightened as follows:

$$\sum_{i \in M: j \in N^i} x_{ij} \leq b_j y_j \quad \forall j \in N. \quad (14)$$

Furthermore, across 10,000 instances, we observe that $a_i < b_j$ for $i \in M, j \in N^i$. As a result, constraints (9) are dominated by (14) and thus eliminated from the model. Now we call (8), (11) – (14) as the “tightened model”.

In our implementation, we first apply the rounding heuristic method to the root LP solution to the original model and then to the optimal LP solution to tightened model. We then use RINS (Danna et al., 2005) to further improve the incumbent. In particular, we define and solve a sub-MIP based on the tightened model, using its LP solution and the current incumbent as a guide to fix part of the binary variables.

4.3. Anonymous Problem

Though the concrete MILP formulation is not available, we discover some pattern from anonymous instances in the LP file format. In particular, one can define a planning horizon and associate with each discrete variable a time period. The details can be deduced from the constraint hypergraph (Rossi et al., 2006) in which every node represents a discrete variable and every edge joins a pair of variables if they occur together in a constraint. Let H denote the planning horizon and $h \in H$ denote a time period. We use a heuristic called the “rolling-horizon” method to generate our final high quality primal solution. It consists of the following steps: (i) ignore constraints involving discrete variables with their h values greater than \tilde{H} ; (ii) relax the integrality constraints on variables with their h values greater than \bar{H} ; (iii) fix those discrete variables with their h values less than \hat{H} at the optimal solution from a previous run; (iv) solve the sub-MIP; (v) increase \hat{H} , \bar{H} and \tilde{H} adaptively and then iterate steps (i) - (iv) until $\hat{H} = H$. Before calling the computationally expensive rolling-horizon method, we utilize feasibility pump (Fischetti et al., 2005; Bertacco et al., 2007) to generate the very first solutions and call RENS (Berthold, 2014) once to improve that solution. The RENS model is a sub-MIP defined by fixing those discrete variables with their h values less than $0.9H$ at the incumbent.

5. Solving the Dual Task (Winning Solution by Nuri)

The dual task deals with obtaining tight optimality guarantees (dual bounds) with branching (Khalil et al., 2016; Balcan et al., 2018; Gasse et al., 2019; Gupta et al., 2020; Cappart et al., 2021). Making good branching decisions is regarded as a critical component of modern branch-and-bound solvers. However, it has received little theoretical understanding to this day (Lodi and Zarpellon, 2017). In this task, the environment runs a full-fledged branch-and-cut algorithm with SCIP, and the participants only control the branching decisions of the solver. The metric of interest is the dual integral, which considers the rate at which which the dual bound increases over time. Fruther, all primal heuristics are deactivated to focus only on proving optimality using branching. Execution time was limited to 15 minutes.

We propose Knowledge Inheriting Dataset Aggregation (KIDA), which trains a neural model to decide branching variables with an enhanced version of Dataset Aggregation and a surprisingly effective Model Weight Averaging (MWA) trick. KIDA consists of three steps. First, candidate models are trained by combining the ideas of DAgger (Ross et al., 2011) and Born-Again Neural Networks (Furlanello et al., 2018) to imitate the Strong Branching (Achterberg et al., 2005) heuristics on an expanded dataset. Then, we apply the MWA trick to derive more candidate models. Finally, the model with the highest cumulated reward on the validation set is selected as the final model to be used for testing in the deployment environment. We demonstrate that KIDA achieves top performance on the benchmarks of Balanced Item Placement and Anonymous Problem with a single model, surpassing methods that purely imitate Strong Branching heuristics. **Limitations of Imitating Strong Branching** Classical approaches (Gasse et al., 2019; Nair et al., 2020) to train neural models to decide branching variables rely on imitation of Strong Branching heuristics. However, on the two benchmarks of Balanced Item Placement and Anonymous Problem, we find that having good performance measured by the accuracy of imitating Strong Branching (SBA) does not always lead to good performance on the deployment environment measured as cumulated reward (CR), as shown in Table 1. On the other hand, though CR is a reliable evaluation metric that is highly consistent between validation and

deployment, as shown in Table 1, it is a sparse signal and cannot be used directly in an Imitation Learning framework. In light of these, we propose to first train models to imitate Strong Branch heuristics, and then use a greedy search to select the final model from trained models and their derived weight averaging models. We also note on the benchmark of Workload Apportionment, SBA almost completely fails to correlate with the deployment-time CR as models with high SBA are often inferior to simple random strategy, as shown in Table 2.

Table 1: Performance Comparison of Models of Different Epochs on the benchmark of Balanced Item Placement. SBA is not strictly related to CR

Epoch	Top 1 SBA	Top 3 SBA	Cum. Reward on Validation Set	Cum. Reward on Test Set
1	0.780	0.932	5202.6	5028.6
5	0.803	0.948	5545.7	5289.3
10	0.808	0.946	5131.9	4807.9
20	0.810	0.947	5038.8	4840.7

Enhanced Dataset Aggregation We find that Dataset Aggregation (DAgger) (Ross et al., 2011) improves SBA of trained models and helps to get better final performance in the deployment environment. We are also inspired by Born-Again Neural Networks (Furlanello et al., 2018) to utilize all trained models from the whole training process. The training of the current model depends on the data generated with the last model. By the end of this step, we have a collection of candidate models that can be leveraged in the next step.

Model Weight Averaging and Greedy Search In the setting of Dual Task, Model Ensembles (Dietterich, 2000; Kidziński et al., 2018), which simply averages outputs of multiple models, doesn’t work because the increase in run time by using multiple models will lead to fewer rounds of interaction with the underlying SCIP solver, and consequently degrading performance. In our experiments, averaging the weights of models trained in different epochs (Tarvainen and Valpola, 2017) doesn’t have any positive effect either. In contrast, we consider building a new model π_{avg} by averaging the weights of different trained models during the DAgger process. Formally, for Ω models obtained from DAgger with parameters $(\theta_0, \theta_1, \dots, \theta_{\Omega-1})$, the parameters of π_{avg} are obtained by: $\theta_{avg} = \sum_{i=0}^{\Omega-1} \theta_i / \Omega$. In practice, we perform a greedy grid search over different model combinations and different Ω to select the best averaged model with the highest CR. Note original models are also included in the search as the particular case of $\Omega = 1$.

Results

The final performance of our methods on benchmarks is shown in Table 2. Baseline results are strictly from the official code. For KIDA, we select the best-performing model by greedy search over $\Omega = 1, 2, \dots, 5$.

Table 2: The Cumulated Reward in Three Benchmarks of Different Methods

	Random	Baseline	KIDA
Balanced Item Placement	3300.7	4937.8	7561.6 ($\Omega = 3$)
Workload Apportionment	624928.9	624043.6	623996.0
Anonymous Problem	31145708.4	30965031.6	32832618.7 ($\Omega = 1$)

Balanced Item Placement: a KIDA model with $\Omega = 3$ overtakes the other methods with a significant margin.

Workload Apportionment: surprisingly, the random strategy is significantly better than the baseline model in this benchmark. This may be attributed to the failure of Strong Branching heuristics for this set.

Anonymous Problem: the best KIDA model we find is a single DAgger model without averaging. Applying KIDA with $\Omega \geq 2$ leads to a slight decline in deployment-time CR.

6. Solving the Configuration Task (Winning Solution by EI-OROAS)

The configuration task deals with deciding on a good parameterization of the solver for a given problem instance (Hutter et al., 2011). The environment required for this task is more straightforward than for the two previous ones since it involves only a single decision for the agents, i.e., contextual bandit problem). Participants were allowed to tune any of the existing parameters of SCIP. They could choose between providing a fixed set of parameters that work well on average for each problem benchmark or producing instance-specific parameterizations based on the characteristics of each instance. The metric of interest for this task was the primal-dual gap integral, which combines both improvements from the dual and from the primal side over time. To compute this metric unambiguously, even when no primal or dual bound exists, an initial primal bound (trivial solution value) and an initial dual bound (pre-computed root LP solution value) for each instance are both provided. Execution time was limited to 15 minutes.

The configuration task deals with deciding on a good parameterization of the solver for a given problem instance. Specifically, there are two questions need to answer: 1. How to search good parameters for seen instances? 2. How to recommend parameters for unseen instances? With a basic exploration, we conclude the following challenges: 1) Numerous, heterogeneous and conditional dependent configurable parameters 2) Expensive optimization cost for open instances and Limited samples for anonymous set. In our solution, we proposed a novel space reduction pipeline and adopted HEBO (Cowen-Rivers et al., 2020) as optimizer to search good parameter, while a ML-based classifier to recommend parameters.

Implementation of method First, the search space is reduced to accommodate Bayesian optimization. Expert experiences were introduced to filter insignificant parameters. Then, fitting a XGboost observer to tighten the region of candidate parameters again with Gini coefficients. The search space Ω will be split to k sub-space (om)at the start, then optimized one by one with broad first search order. Algorithm 1 presents the pseudo-code for search. Specifically, a two-layer loop is used to solve the problem. The outer loop optimizes each sub-space, and the inner loop is a standard BO method that minimizes the objective function. Some comments are all as follows. 1) *Full solution initialization:* The x^* is assigned the SCIP default value for initialization and replaced by optimal partial solution x^i ($i = 1, \dots, k$) until the end of the iterations. 2) *Partial solution initialization:* The sub-space Ω_i is defined as the current search space, and the initial dataset \mathcal{D}_0^i is sampled in this space. 3) *Model construction and optimization:* The surrogate model construction and acquisition function optimization is following the HEBO. 4) *Objective evaluation:* The new suggested solutions $x_{1:q}^i$ will supplemented by x^* into full solutions, configuring for SCIP, and solving for all instances in \mathcal{I} . the mean value of the Primal-dual gap integral for all instances is returned as an objective value. 5) *Update full solution:* When the inner loop ends, the optimal solutions x_i will update the global full optimal solution x^* .

Empirical study We solve the problem as described above, where the anonymous dataset is divided into three clusters according to the features. The statistical results of SCIP solver with

Table 3: The statistical results of primal-dual gap by default and tuned parameters of SCIP.

Dataset	Default Performance	Tuned Performance	Improvement
Item Placement	16942	8781	×1.92
Load Balancing	22168	9499	×2.33
Anonymous-c1	277e+4	257e+4	×1.08
Anonymous-c2	1133e+6	730e+6	×1.55
Anonymous-c3	726e+4	574e+4	×1.26

Algorithm 1: Pseudocode of Parameter Adaptive Bayesian Optimization

Input : $\Omega = \{\Omega_1, \Omega_2, \dots, \Omega_k\}$: Search space; \mathcal{I} : Instances set; N : Total number of iterations of each sub-search space

$x^* \leftarrow$ SCIP default value in Ω . /* full solution initialization */

for $i \leftarrow 1$ **to** k **do**

Initialize \mathcal{D}_0^i by random sample from sub-search space Ω_i . /* partial solution initialization */

for $j \leftarrow 0$ **to** $N - 1$ **do**

Fit a surrogate model to current dataset \mathcal{D}_j^i . /* modeling */

Find q solutions $x_{1:q}^i$ by maximizing three acquisition functions. /* suggest */

Evaluate new partial solutions $(x_{1:q}^i)$ by querying the function to get $y_{1:q} = f(x_{1:q}^i, x^*, \mathcal{I})$. /* evaluate */

Update the dataset creating $\mathcal{D}_{j+1}^i = \mathcal{D}_j^i \cup \{x_l^i, y_l\}_{l=1}^q$. /* tell */

end

$x^* \leftarrow x^* \cup \operatorname{argmin}_{x^i \in \mathcal{D}^i}$. /* update full optimum solution */

end

Output : x^* (the best-performing parameters set for SCIP solver on \mathcal{I}).

default and tuned parameters are summarize in Table 3. The performance of the SCIP has been improved, with a maximum of 2.33 times improvement on Load Balancing dataset and minimum of 1.08 times on Anonymous-c1. The experimental results demonstrate the effectiveness and robustness of our method in the MIP solver configuration.

7. Conclusion

The goal of this competition was to foster the design of innovative methods to improve state-of-the-art combinatorial optimization solvers by replacing key heuristic components with machine learning models. To that, we proposed three challenging benchmarks. In total, we received 12 for the primal task 23 for the dual task, and 15 for the configuration task. We provide descriptions of some of the participants' solutions on the competition website.² The results indicate that machine learning for combinatorial optimization has potential, although more work must be done before it becomes relevant for practical, real-world use. We plan to maintain this competition across the years to monitor performance improvements over the years.

Acknowledgments

The event was sponsored by the Artificial Intelligence Journal, as well as Compute Canada, Calcul Québec, and Westgrid who graciously provided the compute resources and the prize money. We finally thank all the participants.

2. <https://www.ecole.ai/2021/ml4co-competition/>

References

- Tobias Achterberg, Thorsten Koch, and Alexander Martin. Branching rules revisited. *Operations Research Letters*, 33(1):42–54, 2005.
- Maria-Florina Balcan, Travis Dick, Tuomas Sandholm, and Ellen Vitercik. Learning to branch. In *International Conference on Machine Learning*, pages 344–353. PMLR, 2018.
- Livio Bertacco, Matteo Fischetti, and Andrea Lodi. A feasibility pump heuristic for general mixed-integer problems. *Discrete Optimization*, 4(1):63–76, 2007.
- Timo Berthold. Rens. *Mathematical Programming Computation*, 6(1):33–54, 2014.
- Quentin Cappart, Thierry Moisan, Louis-Martin Rousseau, Isabeau Prémont-Schwarz, and Andre A Cire. Combining reinforcement learning and constraint programming for combinatorial optimization. In *AAAI Conference on Artificial Intelligence*, volume 35, pages 3677–3687, 2021.
- Alexander I Cowen-Rivers, Wenlong Lyu, Rasul Tutunov, Zhi Wang, Antoine Grosnit, Ryan Rhys Griffiths, Hao Jianye, Jun Wang, and Haitham Bou Ammar. An empirical study of assumptions in bayesian optimisation. *arXiv preprint arXiv:2012.03826*, 2020.
- Emilie Danna, Edward Rothberg, and Claude Le Pape. Exploring relaxation induced neighborhoods to improve mip solutions. *Mathematical Programming*, 102(1):71–90, 2005.
- Thomas G Dietterich. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, pages 1–15. Springer, 2000.
- Matteo Fischetti, Fred Glover, and Andrea Lodi. The feasibility pump. *Mathematical Programming*, 104(1):91–104, 2005.
- Tommaso Furlanello, Zachary Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar. Born again neural networks. In *International Conference on Machine Learning*, pages 1607–1616. PMLR, 2018.
- Gerald Gamrath, Daniel Anderson, Ksenia Bestuzheva, Wei-Kun Chen, Leon Eifler, Maxime Gasse, Patrick Gemander, Ambros Gleixner, Leona Gottwald, Katrin Halbig, Gregor Hendel, Christopher Hojny, Thorsten Koch, Pierre Le Bodic, Stephen J. Maher, Frederic Matter, Matthias Miltenberger, Erik Mühmer, Benjamin Müller, Marc E. Pfetsch, Franziska Schlösser, Felipe Serrano, Yuji Shinano, Christine Tawfik, Stefan Vigerske, Fabian Wegscheider, Dieter Weninger, and Jakob Witzig. The SCIP Optimization Suite 7.0. Technical report, Optimization Online, March 2020. URL http://www.optimization-online.org/DB_HTML/2020/03/7705.html.
- Maxime Gasse, Didier Chételat, Nicola Ferroni, Laurent Charlin, and Andrea Lodi. Exact combinatorial optimization with graph convolutional neural networks. In *Advances in Neural Information Processing Systems*, 2019.
- Prateek Gupta, Maxime Gasse, Elias Khalil, Pawan Mudigonda, Andrea Lodi, and Yoshua Bengio. Hybrid models for learning to branch. 2020.

- Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. Sequential model-based optimization for general algorithm configuration. In *International Conference on Learning and Intelligent Optimization*, pages 507–523. Springer, 2011.
- Elias Khalil, Pierre Le Bodic, Le Song, George Nemhauser, and Bistra Dilkina. Learning to branch in mixed integer programming. In *AAAI Conference on Artificial Intelligence*, volume 30, 2016.
- Elias Khalil, Hanjun Dai, Yuyu Zhang, Bistra Dilkina, and Le Song. Learning combinatorial optimization algorithms over graphs. In *Advances in Neural Information Processing Systems*, pages 6348–6358, 2017.
- Lukasz Kidziński, Sharada Prasanna Mohanty, Carmichael F Ong, Zhewei Huang, Shuchang Zhou, Anton Pechenko, Adam Stelmazczyk, Piotr Jarosik, Mikhail Pavlov, Sergey Kolesnikov, et al. Learning to run challenge solutions: Adapting reinforcement learning methods for neuromusculoskeletal environments. In *The NIPS’17 Competition: Building Intelligent Systems*, pages 121–153. Springer, 2018.
- Andrea Lodi and Giulia Zarpellon. On learning and branching: a survey. *Top*, 25(2):207–236, 2017.
- Vinod Nair, Sergey Bartunov, Felix Gimeno, Ingrid von Glehn, Pawel Lichocki, Ivan Lobov, Brendan O’Donoghue, Nicolas Sonnerat, Christian Tjandraatmadja, Pengming Wang, et al. Solving mixed integer programs using neural networks. *arXiv preprint arXiv:2012.13349*, 2020.
- Mohammadreza Nazari, Afshin Oroojlooy, Martin Takáč, and Lawrence V Snyder. Reinforcement learning for solving the vehicle routing problem. In *Advances in Neural Information Processing Systems*, pages 9861–9871, 2018.
- Dimitri J Papageorgiou, George L Nemhauser, Joel Sokol, Myun-Seok Cheon, and Ahmet B Keha. Mirplib—a library of maritime inventory routing problem instances: Survey, core model, and benchmark results. *European Journal of Operational Research*, 235(2):350–366, 2014.
- Antoine Prouvost, Justin Dumouchelle, Lara Scavuzzo, Maxime Gasse, Didier Chételat, and Andrea Lodi. Ecole: A gym-like library for machine learning in combinatorial optimization solvers. *arXiv preprint arXiv:2011.06069*, 2020.
- Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011.
- Francesca Rossi, Peter Van Beek, and Toby Walsh. *Handbook of constraint programming*. Elsevier, 2006.
- Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *arXiv preprint arXiv:1703.01780*, 2017.