Ecole: A Library for Learning Inside MILP Solvers

Antoine Prouvost[1,2], Justin Dumouchelle[1], Maxime Gasse[1,2], Didier Chételat[1], and Andrea Lodi[1,2]

`{firstname}.{lastname}@polymtl.ca`

[1]Canada Excellence Research Chair in Data Science for Decision Making, École Polytechnique de Montréal
[2]Mila, Quebec Artificial Intelligence Institute

Abstract

In this paper we describe Ecole (Extensible Combinatorial Optimization Learning Environments), a library to facilitate integration of machine learning in combinatorial optimization solvers. It exposes sequential decision making that must be performed in the process of solving as Markov decision processes. This means that, rather than trying to predict solutions to combinatorial optimization problems directly, Ecole allows machine learning to work in cooperation with a state-of-the-art a mixed-integer linear programming solver that acts as a controllable algorithm. Ecole provides a collection of computationally efficient, ready to use learning environments, which are also easy to extend to define novel training tasks. Documentation and code can be found at `https://www.ecole.ai`.

1   Introduction

Combinatorial optimization algorithms play a crucial role in our societies, for tackling a wide range of decision problems arising in, but not limited to, transportation, supply chain, energy, finance and scheduling [Paschos, 2013]. These optimization problems, framed as mathematical programs, are inherently hard to solve, forcing practitioners to constantly develop and improve existing algorithms. As a result, general-purpose mathematical solvers typically rely on a large number of handcrafted heuristics that are critical to efficient problem solving, but whose interplay is usually not well understood and is exponentially hard to analyse. These heuristics can be regarded as having been *learned* by human experts through trial and error, on public (or private) data-sets of problems such as MIPLIB [Gleixner et al., 2021].

At the same time, traditional solvers typically disregard the fact that some applications require to solve similar problems repeatedly, and tackle each new problem independently, without leveraging any knowledge from the past. In this context, applying machine learning (ML) to combinatorial optimization (CO) appears as a natural idea, and has actually been a topic of interest for quite some time [Smith, 1999]. With the recent success of ML, especially the deep learning (DL) sub-field, there is renewed

appeal to replace some of the heuristic rules inside traditional solvers by statistical models learned from data. The result would be a solver whose performance could be automatically tailored to a given distribution of mathematical optimization problems, which could be either application-specific or general-purpose ones. The reader is referred to Bengio et al. [2020] for a detailed survey on the topic.

In this article, we present Ecole, an open-source library aimed at facilitating the development of ML approaches within general-purpose mixed-integer linear programming (MILP) solvers based on the branch-and-bound (B&B) algorithm. The remainder of this article is organized as follows. In Section 2, we detail the challenges faced by practitioners for applying ML inside combinatorial optimization (CO) solvers. In Section 3, we present existing software that also aim at facilitating the development of ML solutions for CO. In Section 4, we provide background on mixed-integer linear programming and the branch-and-bound algorithm, as well as the concepts of Markov decision process and reinforcement learning. In Section 5, we present our formulation of control problems arising in mathematical solvers as Markov decision processes, and in Section 6, we showcase the Ecole interface and how it relates to this formulation. In Section 7, we compare the computing performance of Ecole for extracting solver features, compared to existing implementations from the literature. Finally, we conclude with a discussion on future plans for Ecole in Section 8.

2 Motivation

Building the appropriate software to apply ML inside of a B&B solver is not an easy task, and requires a deep knowledge of the solver. It may take months of software engineering before researchers can focus on the actual ML algorithm, and the engineering endeavors can be dissuasive. For example, it suffices to look at research articles with public software implementation [He et al., 2014, Gasse et al., 2019, Gupta et al., 2020, Zarpellon et al., 2020] to get an idea of the complexity of the required code base. Not to mention the fact that such implementations can themselves contain bugs, and will quickly become outdated.

Solvers such as SCIP [Gamrath et al., 2020], CPLEX [IBM, 2020], and Gurobi [Gurobi Optimization LLC, 2020], expose their application programming interface (API) in the c programming language, while the state-of-the-art tools for ML such as Scikit-Learn [Pedregosa et al., 2011], PyTorch [Paszke et al., 2017], and TensorFlow [Abadi et al., 2016] exist primarily in Python. Advanced software engineering skills are necessary to interface both ecosystems, and the room for errors is large, especially if additional time is not invested to write tests for the code. Once these hardships are overcome, the resulting implementation may still be slow and lack advanced features such as parallelization (in particular due to the Python global interpreter lock (GIL) that prevents multi-threaded code executions). Furthermore, research software written for particular projects is often difficult to reuse without copy-editing code, as they lack extensible concept abstractions, proper software packaging, and code maintenance.

Ecole is a free and open-source library built around the SCIP solver to address the

aforementioned issues. Several decision problems of interest that arise inside the solver are exposed through an extensible interface akin to OpenAI Gym library [Brockman et al., 2016], a library familiar to ML practitionners. Going further, Ecole aims at improving the reproducibility of scientific research in the area with unified problem benchmarks and metrics, and provides strong default options for new researchers to start with, without the need for an expert knowledge of the inner workings of a mathematical solvers.

3   Related Work

Other libraries have been introduced recently to facilitate the application of ML to operations research. MipLearn [Xavier and Qiu, 2020] is basically aimed at the same goals as Ecole, with a strong focus on customization and extensibility. It supports two competitive commercial solvers, namely CPLEX and Gurobi, but as a result is limited in the type of interactions it offers, and only allows for using ML for solver configuration. In contrast, Ecole only supports the open-source solver SCIP, but allows for repeated decision making, such the selection of branching variables during B&B, which is a cornerstone of the algorithm. ORGym [Hubbs et al., 2020] and OpenGraphGym [Zheng et al., 2020] also offer Gym-like learning environments, for general operations research (OR) problems and for graph-based problems, respectively. Both are aimed at using ML to produce feasible solutions directly, without the need for an MILP solver. As such they do not allow for the exact solving of CO problems. Ecole, on the other hand, benefits from the inherent mathematical guarantees of a mathematical solver, which include the possibility of exact solving. As such Ecole does not necessarily offer a replacement to the existing software in the ML for OR ecosystem, but rather a (nice) complement that fills some existing gaps. For instance, practitioners can use one of the problem benchmarks from MipLearn, ORGym or OpenGraphGym, to generate a collection of instances in a standard format, and then use Ecole for learning to branch via ML.

4   Background

We now introduce formally some key concepts that are relevant for describing Ecole, related to both combinatorial optimization and reinforcement learning.

4.1   Combinatorial Optimization

Mathematical optimization can be used to model a variety of problems. Variables model the decisions to be made, constraints represent physical or structural restrictions, while the objective function defines a measure of cost to be minimized. When the objective and constraints are linear, and some variables are restricted to be integer, the problem is a mixed-integer linear programming (MILP) problem. MILP is an important class of

decision problems but MILP problems are, in general, $\mathcal{NP}$-hard to solve. The B&B algorithm [Land and Doig, 1960] is an implicit enumeration scheme that is generally at the core of the mathematical solvers designed to tackle these problems. The algorithm starts by computing the (continuous) linear programming (LP) relaxation, typically using the Simplex algorithm [Dantzig, 1990]. If the solution respects the integrality constraints, then the solution is optimal. Otherwise, the feasible space is split by branching on (*i.e.*, partitioning the domain of) an integer variable in a way that excludes the solution to the current LP relaxation. The algorithm is then recursively applied to the sub-domain. If a sub-domain LP relaxation is infeasible, or if its objective value does not improve on the best feasible solution, then the algorithm can stop exploring it.

General-purpose solvers, such as SCIP [Gamrath et al., 2020], CPLEX [IBM, 2020], and Gurobi [Gurobi Optimization LLC, 2020], have emerged to provide an enhanced version of the B&B algorithm. They include additional techniques such as presolving [Achterberg et al., 2020], cutting planes [Gomory, 2010, Balas et al., 1993], and primal heuristics [Fischetti and Lodi, 2011], which together with B&B have contributed to drastically reduce the solving time of MILP in the last decades [Bixby and Rothberg, 2007].

4.2 Reinforcement Learning

In reinforcement learning (RL), an agent interacts with an environment and receive rewards as a (indirect) consequence of its actions. The frameworks is defined by Markov decision process (MDP) problems as follows. At every time step t, the agent is in a given state S_t and decides on an action A_t. The agent decisions are modeled by a probabilistic policy function: for a given state s that the agent is in, the agent takes an action a with probability $\pi(a|s)$. As a result, and depending on the unknown dynamics of the environment, the agent transitions into a new state S_{t+1} and receives a reward R_{t+1}. If the agent is in a state s, and takes an action a, then the probability to transition into a new state s' and receiving a reward r is denoted by $\mathbb{P}(s', r|a, s)$. This is illustrated in Figure 1. The process ends when reaching a state defined to be terminal, where the set of possible actions is empty. The objective for the agent is to maximize the expected sum of future rewards.
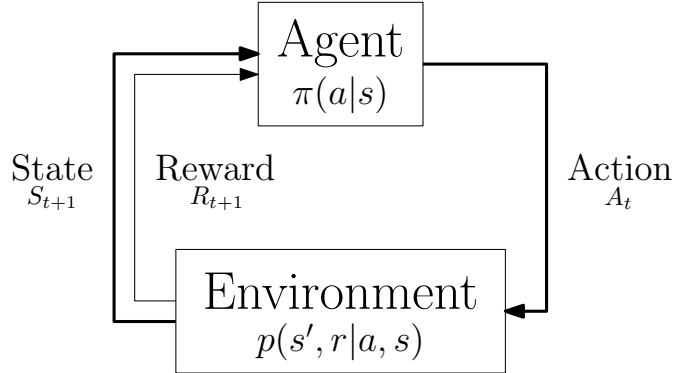
Figure 1: The MDP associated with reinforcement learning, modified from Sutton and Barto [2018].

A sequence of actions and transitions τ is called an *episode*, or *trajectory*. The probability of a given trajectory is given by

$$\mathbb{P}(\tau) = \mathbb{P}(S_0) \prod_t \pi(A_t|S_t)\mathbb{P}(S_{t+1}, R_{t+1}|A_t, S_t). \qquad (1)$$

Common RL algorithms fall in two categories. Policy methods try to directly learn an approximation of the agent policy function π. Value based methods estimate the state-action value function $Q(s, a)$ defined as the expected sum of future rewards given that the agent is in a state s and takes an action a. The associated policy is then computed through approximate dynamic programming [Bertsekas, 2008]. Finally, modern approaches such as actor-critic methods both train a policy (the actor) and an estimate of the value function of the policy (the critic) in a single procedure.

4.3   Imitation Learning

An alternative to find good policies in a MDP is to learn to imitate another expert policy, an approach usually referred to as imitation learning. In this scenario, something usually prevents using the expert directly: for example, the expert might be a human, or might be an expensive algorithm that takes excellent decisions at high computational cost.

The simplest possible approach for imitation learning is behavioral cloning, where state-action pairs are collected by running the expert for a while, and a ML model is trained by imitation learning to predict actions that would be taken by the expert in those states. A downside of this approach is the state distribution mismatch problem: as the student will usually make mistakes, it will deviate from the kinds of states likely to be encountered by the expert and will soon end up in new states much unlike those seen during training. So called active imitation learning methods, such as DAGGER [Ross et al., 2011], try to correct this mismatch to improve final policy performance.

5 Solver Control Problem

Although MILP solvers are usually deterministic processes, in Ecole we adopt a general MDP formulation, introduced in Section 4.2, which allows for non-deterministic processes.

To better map the Ecole interface presented in the next section, we further refine probability (1). First, we split the initial state distribution $\mathbb{P}(S_0)$ using Bayes' rule to introduce the probability distribution of the problem instance I being solved, which can be formulated as $\mathbb{P}(S_0) = \mathbb{P}(I)\mathbb{P}(S_0|I)$. Second, we borrow from the partially-observable MDP framework [Spaan, 2012] and introduce an observation function $\mathcal{O}$ of the state that is used by the agent to make decisions, using all of the past trajectory, that is $\pi(A_t|\mathcal{O}(S_0), ..., \mathcal{O}(S_t))$. Similarly, we reformulate rewards as a function $\mathcal{R}$ of the state and leave it out of the transition probability $\mathbb{P}(S_{t+1}|A_t, S_t)$. This gives us the following final probability distribution of a trajectory:

$$\mathbb{P}(\tau) = \mathbb{P}(I)\mathbb{P}(S_0|I) \prod_t \pi(A_t|H_t)\, \mathbb{P}(S_{t+1}|A_t, S_t), \tag{2}$$

where the history H_t is given by

$$H_t = \{\mathcal{O}(S_0), \mathcal{R}(S_0), A_0, ..., \mathcal{O}(S_{t-1}), \mathcal{R}(S_{t-1}), A_{t-1}, \mathcal{O}(S_t)\}. \tag{3}$$

Ecole provides an API to define environments inside the solver. It lets users sample from those environments according to (2), as well as extract rewards and observations using the functions $\mathcal{R}$ and $\mathcal{O}$. Two environments are currently implemented in Ecole.

Configuring The first environment expose the task of algorithm configuration [Hoos, 2011]. The goal is to find well performing SCIP parameters, then let the solver run its course without intervention. This task is akin to what is learned in Hutter et al. [2010]. In this scenario, finding optimal parameters can be framed as a contextual bandit problem[1]. Contextual bandit problems are special cases of RL where the underlying MDP has unit episode length, so it fits naturally in the framework adopted by Ecole as a special case. In this environment, the one and only action to be taken is a mapping of SCIP parameters and associated values.

Branching The second environment implemented allows users to select the variable to branch on each B&B node, as used in Gasse et al. [2019]. The state S_t is defined as the state of the solver on the t^{th} node, and is equivalent to the branching rule callback available in SCIP. In Gasse et al. [2019], the observation function $\mathcal{O}$ is extracting a bipartite graph representation of the solver, and π is a graph neural network (GNN).

Plans for future environments are suggested in Section 8.

---

[1] A good reference on the topic is Lattimore and Szepesvári [2020, Part V].

6 Design Decisions

The API of Ecole is designed for ease of use and extensibility. In this section, we detail some of the key features of Ecole.

6.1 Environment Interface

Environments The interface for using an environment is inspired by the OpenAI Gym [Brockman et al., 2016] library. The main abstraction is the `Environment` class, which is used to encapsulate any control problem, as formulated in Section 5. The Listing 1 provides an example of using Ecole with the `Branching` environment for B&B variable selection. The inner `while` loop spans over a full episode, while the outer `for` loop repeats it multiple times for different problem instances. An episode, starts with a call to `reset`. The method takes as parameter an MILP problem instance from which an initial state will be sampled, and returns an observation of that state, and a Boolean flag indicating whether that state is terminal. Transitions are performed by calling the `step` method, with the action provided by the user (the `policy` function). It returns the observation of the new state, the reward, the flag indicating whether the new state is terminal, and a dictionary of additional non-essential information about the transition.

```python
import ecole

env = ecole.environment.Branching(
    reward_function=ecole.reward.LpIterations() ** 2,
    observation_function=ecole.observation.NodeBipartite(),
)
instances = ecole.instance.IndependentSetGenerator(n_nodes=100)

for _ in range(10):
    obs, action_set, reward_offset, done, info = env.reset(next(instances))
    while not done:
        action = policy(obs, action_set)
        obs, action_set, reward, done, info = env.step(action)
```

Listing 1: Default usage of environments in Python.

Reward and Observation Functions Furthermore, Listing 1 shows how the constructor of `Environment` can be used to specify the rewards and observations to compute. Some observation functions from the literature are provided Ecole, such as the ones used in Gasse et al. [2019] and Khalil et al. [2016]. The listing also demonstrate how new reward functions can be dynamically created by applying mathematical operations (`+`, `-`, `*`, `/`, `**`, `.exp()`...) on them.

Instance Generators Learning inside a solver may require large amount of training instances. Although industry applications can provide extensive datasets of instances of interest, it is also valuable to have on hand generators with good defaults to quickly

experiment ideas. Out of convenience, we provide four families of generators for users to learn from. These are the same families of instances that were used to benchmark the imitation learning method of Gasse et al. [2019].

- Set covering MILP problems generated following [Balas and Ho, 1980];

- Combinatorial auction MILP problems generated following [Leyton-Brown et al., 2000, Section 4.3];

- Capacitated facility location MILP problems generated following [Cornuéjols et al., 1991];

- Independent set MILP problems generated following the procedure of [Bergman et al., 2016, Section 4.6.4] with both Erdos-Renyi [Erdos and Renyi, 1959] and Barabasi-Albert [Barabási and Albert, 1999] graphs.

Table 1 summarizes the key abstractions in Ecole and how they map to the mathematical formulation presented in Section 5. Some elements are further explained in the next section.

Observation function	$\mathcal{O}$	`NodeBipartite()`
Reward function	$\mathcal{R}$	`LpIterations() ** 2`
Instance dist.	$\mathbb{P}(I)$	`IndependentSetGenerator(n_nodes=100)`
State	S_t	`Model`
Cond. initial state dist.	$\mathbb{P}(S_0|I)$	`BranchingDynamics.reset_dynamics`
Policy	$\pi(A_t|H_t)$	`policy`
Transition dist.	$\mathbb{P}(S_{t+1}|A_t, S_t)$	`BranchingDynamics.step_dynamics`

Table 1: Comparing Ecole API to its mathematical formulation.

6.2 Extensibility

OpenAI Gym is designed as a set of benchmarks for ML practitioners. However, in Ecole the tasks also have industry applications, and therefore it was an essential design principle to allow users flexibility in designing the environments. For example, although good defaults bring value, if users decide to train an agent with a customized observations or reward function and this leads in the end to better solving times, this is a net gain and Ecole should allow for it. Thus, environments were made to be customizable, unlike in OpenAI Gym.

In this section, we explain how users can customize the reward function $\mathcal{R}$, the observation function $\mathcal{O}$, the instance distribution $\mathbb{P}(I)$, and the transition dynamics $\mathbb{P}(S_0|I)$ and $\mathbb{P}(S_{t+1}|A_t, S_t)$.

Reward and Observation Functions   Users can create observation or reward functions by creating a class with two methods, as shown in Listing 2. They have access to the state of the MDP, *i.e.*, the underlying SCIP solver, through the `ecole.scip.Model`,

or equivalently a `pyscipopt.Model` object (both being wrappers of a `SCIP*` pointer in `C`). There is no limitation to what an observation can be because they are an abstraction used exclusively by the user.

```python
class Observation:
    ...

class ObservationFunction:
    def before_reset(model: ecole.scip.Model) -> Observation:
        ...

    def extract(model: ecole.scip.Model) -> Observation:
        ...
```

Listing 2: Python interface of an observation function.

Transitions and Initial State   The initial state and transition probability distribution can be customized by creating a `EnvironmentDynamics` object. Its API is similar to that of the `Environment`, with the exception that the two methods `reset_dynamics` and `step_dynamics` solely manipulate the solver, without computing either observations or rewards. Environments are actually wrappers around dynamics that also call the reward and observation functions.

Instance Generators   Instance generators are regular Python generators that output an `ecole.scip.Model`. The users are free to define any new generators that can create problem instances, or read them from file.

For all the abstractions above, existing components of Ecole can easily be reused through composition or inheritance to speed up development.

6.3   Comparison with OpenAI Gym

One objective of developing Ecole is to provide an interface close to the popular OpenAI Gym library. Nonetheless, some differences exist.

Condition Initial State distribution   The `reset` method in Gym does not accept any parameters. In Ecole, it makes little sense to solve the same MILP instance over and over, hence conditioning the initial state probability distribution on the instance is mandatory.

Initial Terminal States   In the Gym interface, initial states cannot be terminal. As a result the `reset` method does not return the boolean termination flag. However, terminal initial states do arise in the environments defined in Ecole. For example, in B&B variable selection, the instance could be solved through preprocessing and never require any branching.

9

Reward Offsets   In Ecole, rewards are also returned on `reset`. This is because rewards usually come from differences of metrics, but users cannot compute the total metric without knowing how the metric evolved until the first decision point. The reset reward is this missing information. For example, the solving time reward reports how much time is spent between each decision and the next state. Summing these rewards give the total time spent solving since the agent was first asked for a decision, but it does not take into account time spent before that point. In B&B variable selection, this time would include preprocessing, and root node cutting plane calculation and LP solving. This pre-decision time is what is returned by reset, allowing it to be summed to the rest of the rewards to find the total solving time.

Action Spaces   In OpenAI Gym, `action_space` is a property of the environment class, that is, it does not change between episodes. In Ecole, not only the set of actions can change from one instance to the next, it can also change between transitions. For instance, in B&B variable selection, the set of valid branching candidates changes on every node.

6.4   Ecosystem

The largest part of Ecole is written in `c++`, and exposes bindings to Python through PyBind11 [Jakob et al., 2017]. Xtensor [Mabille et al., 2016] is used for high level multi-dimensional arrays, and NumPy [Harris et al., 2020] is used for binding them to Python.

Writing Ecole in `c++` provides solid ground for a computationally efficient library that can be made available in multiple programming languages [Vollprecht, 2018].

7   Performance Experiments

Speed is a core principle of Ecole. Besides having most of its codebase in `c++`, selecting carefully the third-party libraries on which it relies (see Section 6.4) and tuning compilation options, other optimizations were leveraged to further speedup Ecole. When using Ecole from Python, calls into the `c++` code are made without copying parameters and return values unless necessary. Furthermore, any function doing significant work, such as data extraction or calls to the SCIP solver, are performed without holding the GIL, allowing Python users to use Python threading capabilities and avoid the overhead that comes with multiprocessing parallelization.

We propose two experiments to illustrate the gains provided by these optimizations and benchmark Ecole performance.[2] The benchmarks were run on a server with Linux OpenSUSE Tumbleweed 20210114, with 32GB of RAM, and and Intel i7-6700K (8 cores, 4.0 GHz) CPU. To analyze the results, we used SciPy [Virtanen et al., 2020], Jupyter [Kluyver et al., 2016], Pandas [McKinney, 2010], Matplotlib [Hunter, 2007], and Seaborn [Waskom and the Seaborn development team, 2020].

---

[2]The code of the experiments is available at `https://github.com/ds4dm/ecole-paper`.

7.1 Ecole Overhead Experiment

The first concern was to understand whether using Ecole without extracting any data creates any overhead compared to using SCIP directly. In particular, SCIP offers the possibility to select a branching variable through the use of a callback function, while the Ecole interface presented in Section 6.1 wraps the callback function to select branching variables iteratively (effectively transforming the callback in a stackful coroutine, *i.e.*, a function that can be suspended and resumed).

Using the four instance generators from [Gasse et al., 2019] (implemented in Ecole and presented in Section 6.2), we compare branching on the first available branching candidate using Ecole and vanilla SCIP. To speedup the benchmark, we disable presolving, cutting planes, and limit the number of nodes to 100. Over a total of 4500 instances, a one sample t-test shows that the ratio of the wall times is not significantly different from 1.0 (with p-value $< 10^{-50}$). Thus, we can conclude that for a typical usage, Ecole produces no overhead. This is explained by the fact that any possible overhead time is dwarfed in comparison to the time spent solving the problem.

7.2 Observation Functions Experiment

Our second experiment aims at measuring the sole execution time of two observation functions implemented in Ecole: `NodeBipartite` from Gasse et al. [2019] and `Khalil2016` from Khalil et al. [2016] and implemented in Gasse et al. [2019]. Using the four instance generators from Gasse et al. [2019], with no presolving, no cutting planes, and limiting to 100 nodes, we measure their sole execution time. A total of 439 instances were generated.

The results are given in Table 2 for `Khalil2016`. For `NodeBipartite`, it is possible to cache some information computed at the root node for future use but this is incompatible with cutting planes. The results without and with a cache are given in Tables 3 and 4, respectively.

Generator	Wall time Gasse (s)	Wall time Ecole (s)	Ratio
CFLP 100-100	$(1.067 \pm 0.072) * 10^2$	1.050 ± 0.117	102.2 ± 5.9
CFLP 200-100	$(4.381 \pm 0.641) * 10^2$	2.580 ± 0.420	170.2 ± 8.8
CFLP 400-100	$(2.468 \pm 0.354) * 10^3$	6.151 ± 0.768	400.3 ± 13.7
CAuction 100-500	$(5.682 \pm 0.856) * 10^{-1}$	$(3.752 \pm 0.605) * 10^{-2}$	15.19 ± 0.77
CAuction 200-1000	2.654 ± 0.437	$(1.067 \pm 0.190) * 10^{-1}$	24.93 ± 0.74
CAuction 300-1500	7.207 ± 0.718	$(1.967 \pm 0.209) * 10^{-1}$	36.68 ± 0.98
IndependentSet 500	1.531 ± 0.192	$(3.210 \pm 0.285) * 10^{-1}$	4.755 ± 0.212
IndependentSet 1000	8.209 ± 1.463	1.424 ± 0.146	5.726 ± 0.512
IndependentSet 1500	$(2.226 \pm 0.407) * 10$	3.060 ± 0.284	7.228 ± 0.724
SetCover 500-1000	$(9.962 \pm 6.410) * 10^{-2}$	$(9.261 \pm 6.197) * 10^{-3}$	10.26 ± 2.99
SetCover 1000-1000	$(4.351 \pm 2.812) * 10^{-1}$	$(5.136 \pm 3.403) * 10^{-2}$	8.613 ± 0.540
SetCover 2000-1000	1.210 ± 0.627	$(1.980 \pm 1.135) * 10^{-1}$	6.274 ± 0.555

Table 2: Comparison of execution times for `Khalil2016`

.

Generator	Wall time Gasse (s)	Wall time Ecole (s)	Ratio
CFLP 100-100	$(4.232 \pm 0.070) * 10^{-1}$	$(1.209 \pm 0.038) * 10^{-1}$	3.505 ± 0.113
CFLP 200-100	$(8.017 \pm 0.384) * 10^{-1}$	$(2.815 \pm 0.125) * 10^{-1}$	2.848 ± 0.058
CFLP 400-100	1.557 ± 0.016	$(6.179 \pm 0.101) * 10^{-1}$	2.520 ± 0.037
CAuction 100-500	$(7.376 \pm 0.549) * 10^{-2}$	$(4.074 \pm 0.348) * 10^{-3}$	18.14 ± 0.89
CAuction 200-1000	$(9.797 \pm 0.473) * 10^{-2}$	$(8.753 \pm 0.566) * 10^{-3}$	11.21 ± 0.36
CAuction 300-1500	$(1.160 \pm 0.016) * 10^{-1}$	$(1.356 \pm 0.047) * 10^{-2}$	8.564 ± 0.226
IndependentSet 500	$(6.806 \pm 0.043) * 10^{-1}$	$(1.679 \pm 0.017) * 10^{-1}$	4.053 ± 0.028
IndependentSet 1000	2.564 ± 0.012	$(8.849 \pm 0.090) * 10^{-1}$	2.898 ± 0.026
IndependentSet 1500	6.697 ± 0.421	2.046 ± 0.015	3.273 ± 0.201
SetCover 500-1000	$(4.347 \pm 2.556) * 10^{-2}$	$(3.986 \pm 2.413) * 10^{-3}$	10.25 ± 2.94
SetCover 1000-1000	$(1.015 \pm 0.283) * 10^{-1}$	$(1.238 \pm 0.412) * 10^{-2}$	8.368 ± 0.749
SetCover 2000-1000	$(1.768 \pm 0.304) * 10^{-1}$	$(2.896 \pm 0.628) * 10^{-2}$	6.179 ± 0.469

Table 3: Comparison of execution times for `NodeBipartite` without cache.

Generator	Wall time Gasse (s)	Wall time Ecole (s)	Ratio
CFLP 100-100	$(1.472 \pm 0.019) * 10^{-1}$	$(9.381 \pm 0.287) * 10^{-2}$	1.570 ± 0.047
CFLP 200-100	$(2.891 \pm 0.123) * 10^{-1}$	$(1.958 \pm 0.083) * 10^{-1}$	1.477 ± 0.031
CFLP 400-100	$(6.093 \pm 0.058) * 10^{-1}$	$(3.818 \pm 0.065) * 10^{-1}$	1.596 ± 0.024
CAuction 100-500	$(1.407 \pm 0.095) * 10^{-2}$	$(2.983 \pm 0.224) * 10^{-3}$	4.720 ± 0.098
CAuction 200-1000	$(1.988 \pm 0.081) * 10^{-2}$	$(6.327 \pm 0.308) * 10^{-3}$	3.144 ± 0.050
CAuction 300-1500	$(2.484 \pm 0.021) * 10^{-2}$	$(9.486 \pm 0.152) * 10^{-3}$	2.619 ± 0.030
IndependentSet 500	$(1.371 \pm 0.019) * 10^{-1}$	$(7.254 \pm 0.090) * 10^{-2}$	1.890 ± 0.018
IndependentSet 1000	$(5.268 \pm 0.047) * 10^{-1}$	$(3.117 \pm 0.032) * 10^{-1}$	1.690 ± 0.016
IndependentSet 1500	1.305 ± 0.008	$(7.088 \pm 0.046) * 10^{-1}$	1.842 ± 0.012
SetCover 500-1000	$(9.406 \pm 4.883) * 10^{-3}$	$(2.836 \pm 1.678) * 10^{-3}$	3.435 ± 1.160
SetCover 1000-1000	$(1.851 \pm 0.432) * 10^{-2}$	$(7.153 \pm 2.052) * 10^{-3}$	2.673 ± 0.374
SetCover 2000-1000	$(2.791 \pm 0.399) * 10^{-2}$	$(1.422 \pm 0.267) * 10^{-2}$	1.985 ± 0.145

Table 4: Comparison of execution time for `NodeBipartite` with cache.

As can be seen, Ecole is systematically faster than the literature. The shorter execution times in Ecole can be explained by the fact that the code from [Gasse et al., 2019] was a proof of concept and was not extensively optimized.

8   Conclusions and Future Work

Ecole offers researchers an efficient and well designed interface to the SCIP solver without compromising on customizability.

A current limitation of the library is that only SCIP is supported as a back-end solver. Since commercial, closed-source solvers such as CPLEX [IBM, 2020] and Gurobi [Gurobi Optimization LLC, 2020] are very popular in industry, it would be natural to extend the library to support them as potential back-ends. For now, their closed-source nature limits this possibility, but we hope that interest in the current version of Ecole will lead solver developers to facilitate interfacing with ML libraries in the future.

Future work on Ecole will involve developing new environments, such as node selection and cutting planes selection, as well as new observation and reward functions, such as primal, dual and primal-dual integral metrics. In addition, support for out-of-the-box parallelism would be useful to cope with computationally expensive environments.

Acknowledgments

References

Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, pages 265–283, 2016.

Tobias Achterberg, Robert E. Bixby, Zonghao Gu, Edward Rothberg, and Dieter Weninger. Presolve reductions in mixed integer programming. *INFORMS Journal on Computing*, 32(2):473–506, 2020. doi: 10.1287/ijoc.2018.0857.

Egon Balas and Andrew Ho. Set covering algorithms using cutting planes, heuristics, and subgradient optimization: a computational study. In *Combinatorial Optimization*, pages 37–60. Springer, 1980.

Egon Balas, Sebastián Ceria, and Gérard Cornuéjols. A lift-and-project cutting plane algorithm for mixed 0–1 programs. *Mathematical programming*, 58(1):295–324, 1993.

Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.

Yoshua Bengio, Andrea Lodi, and Antoine Prouvost. Machine learning for combinatorial optimization: a methodological tour d'horizon. *European Journal of Operational Research*, 2020.

David Bergman, Andre A Cire, Willem-Jan Van Hoeve, and John Hooker. *Decision diagrams for optimization*, volume 1. Springer, 2016.

Dimitri P. Bertsekas. Approximate dynamic programming, 2008.

Robert Bixby and Edward Rothberg. Progress in computational mixed integer programming—a look back from the other side of the tipping point. *Annals of Operations Research*, 149(1):37–41, Feb 2007. ISSN 1572-9338. doi: 10.1007/s10479-006-0091-y. URL https://doi.org/10.1007/s10479-006-0091-y.

Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016.

Gérard Cornuéjols, Ranjani Sridharan, and Jean-Michel Thizy. A comparison of heuristics and relaxations for the capacitated plant location problem. *European journal of operational research*, 50(3):280–297, 1991.

George B Dantzig. Origins of the simplex method. In *A history of scientific computing*, pages 141–151. Association for Computing Machinery, 1990.

Paul Erdos and Alfréd Renyi. On random graphs. *Publicationes Mathematicae*, pages 290–297, 1959.

Matteo Fischetti and Andrea Lodi. Heuristics in Mixed Integer Programming. In *Wiley Encyclopedia of Operations Research and Management Science.* Wiley Online Library, 2011. ISBN 9780470400531. doi: 10.1002/9780470400531.eorms0376.

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. ZIB-Report 20-10, Zuse Institute Berlin, 3 2020.

Maxime Gasse, Didier Chetelat, Nicola Ferroni, Laurent Charlin, and Andrea Lodi. Exact combinatorial optimization with graph convolutional neural networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 15580–15592. Curran Associates, Inc., 2019.

Ambros Gleixner, Gregor Hendel, Gerald Gamrath, Tobias Achterberg, Michael Bastubbe, Timo Berthold, Philipp M. Christophel, Kati Jarck, Thorsten Koch, Jeff Linderoth, Marco Lübbecke, Hans D. Mittelmann, Derya Ozyurt, Ted K. Ralphs, Domenico Salvagnin, and Yuji Shinano. MIPLIB 2017: Data-Driven Compilation of the 6th Mixed-Integer Programming Library. *Mathematical Programming Computation*, 2021. doi: 10.1007/s12532-020-00194-3. URL https://doi.org/10.1007/s12532-020-00194-3.

Ralph E Gomory. Outline of an algorithm for integer solutions to linear programs and an algorithm for the mixed integer problem. In *50 Years of Integer Programming 1958-2008*, pages 77–103. Springer, 2010.

Prateek Gupta, Maxime Gasse, Elias B Khalil, M Pawan Kumar, Andrea Lodi, and Yoshua Bengio. Hybrid models for learning to branch. In *Advances in neural information processing systems*, 2020.

Gurobi Optimization LLC. *Gurobi Optimizer Reference Manual*, 2020. URL http://www.gurobi.com.

Charles R Harris, K Jarrod Millman, Stéfan J van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. Array programming with numpy. *Nature*, 585(7825):357–362, 2020.

He He, Hal Daume III, and Jason M Eisner. Learning to search in branch and bound algorithms. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3293–3301. Curran Associates, Inc., 2014.

Holger H Hoos. Automated algorithm configuration and parameter tuning. In *Autonomous search*, pages 37–71. Springer, 2011.

Christian D. Hubbs, Hector D. Perez, Owais Sarwar, Nikolaos V. Sahinidis, Ignacio E. Grossmann, and John M. Wassick. Or-gym: A reinforcement learning library for operations research problems, 2020.

J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science Engineering*, 9(3):90–95, 2007. doi: 10.1109/MCSE.2007.55.

Frank Hutter, Holger H. Hoos, and Kevin Leyton-Brown. Automated configuration of mixed integer programming solvers. In Andrea Lodi, Michela Milano, and Paolo Toth, editors, *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, pages 186–202, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

IBM. *CPLEX Optimizer User Manual*, 2020. URL https://www.ibm.com/analytics/cplex-optimizer.

Wenzel Jakob, Jason Rhinelander, and Dean Moldovan. *pybind11 – Seamless operability between C++11 and Python*, 2017. URL https://github.com/pybind/pybind11.

Elias Khalil, Pierre Le Bodic, Le Song, George Nemhauser, and Bistra Dilkina. Learning to branch in mixed integer programming. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.

Thomas Kluyver, Benjamin Ragan-Kelley, Fernando Pérez, Brian Granger, Matthias Bussonnier, Jonathan Frederic, Kyle Kelley, Jessica Hamrick, Jason Grout, Sylvain Corlay, Paul Ivanov, Damián Avila, Safia Abdalla, Carol Willing, and Jupyter development team. Jupyter notebooks - a publishing format for reproducible computational workflows. In Fernando Loizides and Birgit Scmidt, editors, *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, pages 87–90, Netherlands, 2016. IOS Press. URL https://eprints.soton.ac.uk/403913/.

A. H. Land and A. G. Doig. An Automatic Method of Solving Discrete Programming Problems. *Econometrica*, 28(3):497–520, 1960. ISSN 0012-9682. doi: 10.2307/1910129.

Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.

Kevin Leyton-Brown, Mark Pearson, and Yoav Shoham. Towards a universal test suite for combinatorial auction algorithms. In *Proceedings of the 2nd ACM conference on Electronic commerce*, pages 66–76, 2000.

Johan Mabille, Sylvain Corlay, and Wolf Vollprecht. *xtensor: Multidimensional arrays with broadcasting and lazy computing*, 2016. URL https://github.com/xtensor-stack/xtensor.

Wes McKinney. Data Structures for Statistical Computing in Python. In Stéfan van der Walt and Jarrod Millman, editors, *Proceedings of the 9th Python in Science Conference*, pages 56 – 61, 2010. doi: 10.25080/Majora-92bf1922-00a.

Vangelis Th Paschos. *Applications of combinatorial optimization*. John Wiley & Sons, 2013.

Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NeurIPS 2017 Autodiff Workshop*, 2017. URL `https://openreview.net/forum?id=BJJsrmfCZ`.

Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *Journal of machine learning research*, 12(Oct):2825–2830, 2011.

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.

Kate A. Smith. Neural Networks for Combinatorial Optimization: A Review of More Than a Decade of Research. *INFORMS Journal on Computing*, 11(1):15–34, February 1999. ISSN 1091-9856. doi: 10.1287/ijoc.11.1.15.

Matthijs TJ Spaan. Partially observable markov decision processes. In *Reinforcement Learning*, pages 387–414. Springer, 2012.

Richard S Sutton and Andrew G Barto. *Reinforcement Learning: An Introduction*. MIT press Cambridge, second edition, 2018. ISBN 9780262039246. URL `http://incompleteideas.net/book/the-book-2nd.html`.

Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.

Wolf Vollprecht. The xtensor vision, 2018. URL `https://towardsdatascience.com/the-xtensor-vision-552dd978e9ad`.

Michael Waskom and the Seaborn development team. *mwaskom/seaborn*, September 2020. URL `https://doi.org/10.5281/zenodo.592845`.

Alinson S Xavier and Feng Qiu. MIPLearn, 2020. URL `https://anl-ceeesa.github.io/MIPLearn`.

Giulia Zarpellon, Jason Jo, Andrea Lodi, and Yoshua Bengio. Parameterizing branch-and-bound search trees to learn branching policies. *arXiv preprint arXiv:2002.05120*, 2020.

Weijian Zheng, Dali Wang, and Fengguang Song. OpenGraphGym: A parallel reinforcement learning framework for graph optimization problems. In *International Conference on Computational Science*, pages 439–452. Springer, 2020.