FINDING BACKDOORS TO INTEGER PROGRAMS: A MONTE CARLO TREE SEARCH FRAMEWORK

Elias B. Khalil
Department of Mechanical
& Industrial Engineering
University of Toronto

Pashootan Vaezipoor
Department of Computer Science
University of Toronto

Bistra Dilkina
Department of Computer Science
University of Southern California

October 19, 2021

ABSTRACT

In Mixed Integer Linear Programming (MIP), a (strong) *backdoor* is a "small" subset of an instance's integer variables with the following property: in a branch-and-bound procedure, the instance can be solved to global optimality by branching *only* on the variables in the backdoor. Constructing datasets of pre-computed backdoors for widely used MIP benchmark sets or particular problem families can enable new questions around novel structural properties of a MIP, or explain why a problem that is hard in theory can be solved efficiently in practice. Existing algorithms for finding backdoors rely on sampling candidate variable subsets in various ways, an approach which has demonstrated the existence of backdoors for some instances from MIPLIB2003 and MIPLIB2010. However, these algorithms fall short of consistently succeeding at the task due to an imbalance between exploration and exploitation. We propose `BaMCTS`, a Monte Carlo Tree Search framework for finding backdoors to MIPs. Extensive algorithmic engineering, hybridization with traditional MIP concepts, and close integration with the CPLEX solver have enabled our method to outperform baselines on MIPLIB2017 instances, finding backdoors more frequently and more efficiently.

1 Introduction

Hard discrete optimization problems arise in a very wide range of application domains. While the theoretical computer science approach focuses on the design of algorithms with approximation guarantees, many real problems are not amenable to such analyses. In the artificial intelligence and operations research communities, on the other hand, much of the focus is on deriving algorithms that can efficiently produce high-quality solutions (even if without approximation bounds) and/or optimality guarantees (even when the worst-case running time is exponential). In this work, we focus on the latter class of approaches, in particular the branch-and-bound procedure for the exact solution of Mixed Integer Programming (MIP) problems. MIP solvers are celebrated for their ability to solve problems with hundreds of thousands to millions of variables and constraints, a feat which the theory would suggest to be impossible in a reasonable amount of time. Yet, little is known about *why* realistic instances of NP-Hard MIPs can be solved relatively efficiently in practice.

The existence of *backdoor* sets provides one possible answer to this puzzle of empirical tractability in exact MIP solving. This was first reported by Dilkina et al. [2009] who showed that backdoors of sizes 10–20 did exist for some instances (of MIPLIB2003 [Achterberg et al., 2006]) with many hundreds to thousands of integer variables. The intuition is simple: if one need only branch on a handful of integer variables to solve a MIP instance to global optimality (or declare it as infeasible), then it is not difficult to imagine that a full-fledged MIP solver would be able to solve that instance without much branching. This motivates the design of algorithms for finding backdoors, with the following potential use cases: (i) Discovering heretofore unknown structural properties for a family of MIP instances, e.g., a new variant of an independent set or facility location problem, by data-mining backdoors from a large number of instances; (ii) Constructing datasets of pre-computed backdoors that can then be used to train machine learning models for quickly identifying these crucial sets on similar but unseen instances and branching on them for a quick solving

time; (iii) Directly using the backdoor for branching, assuming that the backdoor in question can be computed quickly in advance; such a use case, if realized, could speed up solving times substantially for instances that are challenging for current branching strategies in state-of-the-art MIP solvers.

How, then, does one find backdoors? We will denote the desired backdoor size by $K \in \mathbb{N}^+$. Dilkina et al. [2009] used two forms of *random sampling*: independent *uniform* sampling of K integer variables, and independent *biased* sampling of K integer variables to favor ones that are more fractional (i.e., with fractional parts closer to 0.5) in the solution of the Linear Programming (LP) relaxation of the MIP. Candidate sets are sampled (in parallel) and used for branching in a MIP solver; if global optimality is attained, then a backdoor has been found. The biased sampling strategy proved much more effective than uniform. While useful as a proof-of-concept, random sampling is a *pure exploration* strategy which cannot use knowledge from previous draws to focus future ones towards more promising candidate sets.

Fischetti and Monaci [2011] propose another strategy that leverages the following basic fact about (bounded) mixed 0–1 problems: the set of extreme points (vertices) of the LP-feasible region includes all integer-feasible points. The branch-and-bound tree that results from branching on a strong backdoor must have leaf nodes whose LPs are either infeasible or have integer-feasible optima. This leads to the following equivalent definition for a strong backdoor: a small set $\mathcal{B}$ of 0–1 variables such that each fractional vertex has at least one of its fractional variables in $\mathcal{B}$. If one could enumerate all (exponentially many) fractional vertices, then a Set Covering Problem (SCP) can be formulated to find a backdoor. In practice, the fractional vertices are collected progressively in batches by executing branch-and-bound using the current candidate backdoor. An SCP is solved at each iteration until all fractional vertices have been "covered" or an early termination criterion is met. This method can be extended heuristically for general MIPs. Fischetti and Monaci [2011] demonstrate that this method's candidate backdoors, which cover many but rarely all fractional vertices, can lead to smaller trees if selected as branching variables early on in the search. However, because (i) there are can be many fractional vertices and (ii) the order in which they are explored is arbitrary (see [Fischetti and Monaci, 2014, Fischetti, 2014] for a discussion), it is unclear if the information that this method uses is actually conducive to backdoors in practice, i.e., it may be *"exploiting" too aggressively*.

Contributions. Thus far, we have argued (i) that finding backdoors is beneficial for a variety of tasks and (ii) that existing algorithms either explore or exploit too much. Towards a more general, effective, and extensible backdoor search algorithm, we propose `BaMCTS` (short for "Backdoor MCTS"), a Monte Carlo Tree Search (MCTS) framework for finding backdoors to MIPs. Our contributions can be summarized as follows:

1. **Backdoor Search as MCTS:** We contribute the first such formulation of the problem. `BaMCTS` can balance exploration and exploitation by design, is conceptually simple and easy to implement, and is extensible in a plug-and-play fashion.

2. **Tight integration with MIP domain knowledge:** To enable a scalable and effective solution, we customize the high-level MCTS procedure to the MIP setting through the use of domain-specific reward functions, action scoring functions and elimination rules, and careful engineering that is tightly coupled with CPLEX, a widely used MIP solver.

3. **Extensive empirical evaluation** reveals that `BaMCTS` vastly outperforms the sampling strategy of Dilkina et al. [2009], finding "better" backdoors in a shorter amount of time on instances from the MIPLIB2017 Benchmark set Gleixner et al. [2021]. We also show that branching on such sets of variables results in smaller search trees or optimality gaps compared to CPLEX with its default branching.

2 Related Work

Backdoors for combinatorial problems The notion of backdoors was first introduced by Williams et al. [2003] for SAT, where it was observed that practical SAT instances often have a small tractable structures. Over the years many approaches were proposed to find backdors in the SAT context Paris et al. [2006], Kottler et al. [2008], Li and Van Beek [2011]. Observing the connection between SAT and MIP, Dilkina et al. [2009] generalized the concept of backdoors from SAT to MIP and proposed random sampling as a method for finding backdoors. Fischetti and Monaci [2011] proposed another strategy for MIP, which to our knowledge remains the state-of-the-art to this day in the MIP context.

MCTS and applications. MCTS has seen a surge of interest thanks to its great success particularly in solving two-player games Silver et al. [2017]. This has led to many attempts in solving combinatorial optimization problems using MCTS by translating the problem into a game. In particular, UCT has been used to guide MIP solvers Sabharwal et al. [2012a] and to solve Quantified Constraint Satisfaction Problems (QCSP) Satomi et al. [2011]. Bertsimas et al. [2014]

applied MCTS to the challenging problem of Dynamic Resource Allocation (DRA) and showed that it can greatly improve problem-specific baselines on large scale instances. Travelling Salesman Problem (TSP), another classical problem, was addressed in Rimmel et al. [2011] via a version of Monte Carlo search with some success. Bandit Search for Constraint Programming (BaSCoP) Loth et al. [2013] applied MCTS to improve the tree-search heuristics of Constraint Programmin (CP) solvers and showed significant improvements on the depth-first search on certain CP benchmarks. More recent works involve successful application of the "neural" MCTS of AlphaZero to solve a variety of NP-Hard graph problems Abe et al. [2019], as well as solving First-Order Logic descriptions of combinatorial problems Xu et al. [2021].

The difference between our setting and a typical combinatorial optimization problem is that our reward function is a black box, as opposed to an analytical objective function; this makes our setting more challenging because evaluations are time-consuming.

3 Technical Background

3.1 Branch-and-Bound for MIP

We are concerned with Mixed Integer Linear Programming (MIP) problems of the form:

$$z^* = \min\{c^T x | Ax \leqslant b, x \in \mathbb{R}^n, x_j \in \mathbb{Z} \, \forall j \in I\},$$

where $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, $c \in \mathbb{R}^n$, and the non-empty set $I \subseteq \{1, ..., n\}$ indexes the integer variables. The vectors in the set $X_{MIP} = \{x \in \mathbb{R}^n | Ax \leqslant b, x_j \in \mathbb{Z} \, \forall j \in I\}$ are *integer-feasible solutions*. An integer-feasible (or simply feasible) solution $x^* \in X_{MIP}$ is *optimal* if $c^T x^* = z^*$.

A MIP can be solved by *branch and bound* [Cook, 2012], an exact algorithm that divides the original MIP into subproblems organized in a binary tree; see Nemhauser and Wolsey [1988] for a textbook exposition. At each node of the tree, an LP relaxation of the sub-problem is solved. If the resulting solution x_N of the LP relaxation at a node N is integral, then it is also a feasible solution to the MIP, i.e. $x_N \in X_{MIP}$. If such an integral solution has an objective value that is better than the best one found so far, it is referred to as the *incumbent*, maintaining that designation until a better solution is found. Otherwise, the node is either pruned, if its lower bound is greater than the incumbent's value, or branched on, resulting in two child nodes that are added to the queue of nodes to be processed.

Pseudocosts are historical quantities, aggregated for each variable during the search, that represent the amount by which a node's LP relaxation value has increased when branching on a given variable. A higher pseudocost score indicates that branching on a variable typically helps make progress towards proving optimality. As such, pseudocosts are at the heart of most typical branching strategies that are used in MIP solvers [Achterberg et al., 2005, Achterberg and Berthold, 2009, Hendel, 2015]. We refer to the appendix for a formal definition.

3.2 Backdoors for MIP

Given a MIP instance defined by the tuple (A, b, c, I), a *strong backdoor* of size $K \lll n$ is a set of integer variables $B, |B| = K, B \subset I$ such that branching exclusively on variables from B results in a provably optimal solution or a proof of infeasibility. We will consider *order-sensitive* strong backdoors, where B is an ordered set: a variable at rank i must be branched on before variables of rank $j > i$. The order in which the backdoor variables are considered for branching can affect the performance of the solver's primal heuristics, which in turn affects pruning in branch and bound. We refer to Dilkina et al. [2009] for a detailed discussion of why order matters in practice in MIP solvers. Throughout the paper, we will use the term "candidate backdoor" to refer to an ordered set of integer variables that is being assessed but that may or may not be a strong backdoor.

3.3 Monte Carlo Tree Search (MCTS)

MCTS is a randomized algorithm for sequential games with a finite horizon and a finite number of actions n [Browne et al., 2012]. At every step of a two-player game, MCTS seeks to identify the next action to take so as to maximize the probability of winning the game. It does so by building out an n-ary tree in which the root node represents the current state of the game, each edge represents a valid action, and a child node represents the extension of its parent's state by playing the action of the corresponding edge. Associated with each terminal state is a scalar reward value.

Because the complete search tree is exponential in size, MCTS is organized into four key steps that can together navigate the exploration-exploitation trade-off in a sensible way, building out only a partial search tree that is sufficient for identifying a good next action. The four steps are:

(a) **Selection:** Given the current search tree, this step deals with *selecting* the nodes in a depth-first dive from the root. Upper Confidence Trees (UCT) [Kocsis and Szepesvári, 2006] is a widely used scoring rule that combines the average observed reward of a node (or state) with a function of the number of visits to the node; nodes with large average rewards and/or small visit counts obtain high UCT scores, balancing exploration and exploitation. The latter is controlled by an appropriately tuned hyperparameter.

(b) **Expansion:** When the selected node has only a subset of its potential children as nodes in the current search tree, one can choose to *expand* the selected node's child set by creating a node for a new action. When the number of possible actions is large, Progressive Widening [Coulom, 2007, Couëtoux et al., 2011] is used to limit the branching factor of the tree and focus on expanding only frequently-visited (likely more promising) nodes. To select a new action to expand with, uniform sampling may be used. A more "exploitative" expansion would select an action that has led to high-reward children in other nodes of the search tree.

(c) **Simulation:** Following the depth-first dive from the root through a sequence of selections (and possibly a final expansion), a non-terminal node may be reached. Because rewards are only observed in terminal states, e.g., when the game concludes and a winner is declared, *simulation* is used to traverse the state space from the node in question to a terminal state, without consolidating this subpath into the search tree. Typically, one picks actions randomly until a terminal state is reached.

(d) **Backpropagation:** Following the simulation, a reward is collected. *Backpropagation* refers to the credit assignment process whereby the reward is passed on to the nodes along the depth-first path in the tree that led to the observed reward. A sum-backup rule is commonplace: each node accumulates rewards every time it is selected. Alternatively, a more aggressive max-backup rule keeps track of the maximum observed reward, see Sabharwal et al. [2012b] for example.

The four-step loop is repeated until a termination condition (e.g., time limit) is reached, following which the action corresponding to the highest-reward or most visited child of the root node is "played". The opponent responds, and the process is repeated starting with a new search tree representing the updated state of the game.

4 **BaMCTS**

Rather than treat MCTS as a black-box algorithm, we instantiate it for the backdoor search setting by incorporating as much domain knowledge about MIP solving as possible.

We are given a MIP instance defined by the tuple (A, b, c, I) and assume a user-defined bound on the backdoor size, $K \in \mathbb{N}^+$; K should be seen as a small constant on the order of 5 to 10, which would imply a branch-and-bound tree with hundreds of nodes.

We view backdoor search as a *single-player*, *deterministic* constraint satisfaction game. The goal is to find an order-sensitive strong backdoor that satisfies the size-K constraint on the backdoor size.

We now define two key elements of MCTS for the backdoor setting. Let $P(I, K)$ denote the set of all permutations of all subset of the integer set I that have size at most K. For example, if $K = 2$ and $I = \{1, 2, 3\}$ then $P(I, K) = \{(), (1), (2), (3), (1, 2), (1, 3), (2, 1), (2, 3), (3, 1), (3, 2)\}$.

- **State:** A state $S \in P(I, K)$ is a permutation of the variables s.t. $|S| \leq K$. A state is terminal iff $|S| = K$. The root of the MCTS tree corresponds to the empty state $()$.

- **Action:** Given a non-terminal state S, a valid action is the index $i \in I$ of an integer variable such that $i \notin S$. Taking action i in state S means that i is appended to the end of the ordered set S, leading to a new state $S' \in P(I, K), |S'| = |S| + 1$ in which the first $|S|$ variables are the same as those in S.

4.1 Building blocks

Candidate evaluation Consider a terminal state $\hat{S}$. To evaluate this candidate, a MIP solver is used to check if it is a backdoor. In particular, we instrument the CPLEX solver to (i) branch only on variables in $\hat{S}$, terminating the branch and bound if either the instance is solved (i.e., a backdoor has been found) or branching is no longer possible for at least one subproblem (a leaf node in the branch-and-bound tree); (ii) respect the ordering implied by $\hat{S}$; (iii) collect auxiliary data such as pseudocosts and search completion information that will be useful for action selection and reward computation.

Reward shaping The goal–finding an order-sensitive strong backdoor–suggests a binary reward function in which a value of 1 is assigned to a terminal state during MCTS iff a backdoor is found, after which the search terminates be-

cause the goal has been achieved. However, this sparse reward structure makes any form of focused search impossible, an issue that we will tackle by using an appropriate reward function which assigns an informative score to a terminal state that is not a backdoor.

A good reward function is one that gives a large, but not maximal, value to a candidate which satisfies most of the conditions that define a strong backdoor. What are those "conditions"? Following candidate evaluation using the MIP solver, a (potentially incomplete) branch-and-bound tree (not to be confused with the MCTS tree) can be observed. For the candidate to be a backdoor, all leaf nodes of the tree must be closed, i.e., they must have LP relaxations that are either infeasible, integer-feasible, or fathomed by bound; these are the conditions that must be simultaneously satisfied. Now consider a candidate which satisfies most but not all leaf node conditions. The *tree weight* [Kilby et al., 2006] is a scoring function that maps a branch-and-bound tree to a value in $[0, 1]$. When all integer variables are binary, the tree weight is defined as the fraction of binary solutions that belong to a closed leaf node; a strong backdoor would have a tree weight of 1 because all leaf nodes are closed and thus all binary solutions are covered. Importantly, the tree weight achieves our desired criterion for a good reward function: it can give meaningful scores to candidate backdoors that are not true strong backdoors but that eliminate many binary solutions.

Consider the (binary) search tree of branch-and-bound for a mixed-binary integer program, and let T_k denote the tree at the k-th iteration, i.e., after k nodes have been expanded. Let F_k denote the subset of tree nodes that are "final" or fathomed due to their LP relaxations being infeasible, mixed-binary feasible, or worse in value than the best integer-feasible solution's value at the time they were expanded. The tree weight assigns to each node $v \in F_k$ a *weight* of $2^{-d(v)}$, where $d(v)$ is the depth of v. The total tree weight of tree T_k then writes:

$$\text{tree-weight}(T_k) = \sum_{v \in F_k} 2^{-d(v)}.$$

At the start of branch and bound, no nodes have been fathomed and so $F_0 = \emptyset$, tree-weight$(T_0) = 0$. This function strictly increases with more fathomed nodes. We refer to section 4.3 of [Hendel et al., 2020] for further details and an illustrative example, and note that other tree search completion metrics therein could potentially be used instead of tree weight.

The role of pseudocosts   Pseudocosts are historical quantities aggregated for each variable during the search. The upwards (downwards) PC of a variable x_j is the average unit objective gain taken over upwards (downwards) branchings on x_j in previous nodes; we refer to this quantity as Ψ_j^+ (Ψ_j^-). Pseudocost branching at node N with LP solution $\check{x}$ consists in computing values:

$$PC_j = \text{score}\Big((\check{x}_j - \lfloor \check{x}_j \rfloor)\Psi_j^-, (\lceil \check{x}_j \rceil - \check{x}_j)\Psi_j^+\Big)$$

and choosing the variable with the largest such value. Typically, the product is used to combine the downwards and upwards values. One standard way to initialize the pseudocost values is by applying strong branching once for each integer variable, at the first node at which it is fractional Linderoth and Savelsbergh [1999]. We will refer to this PC strategy with strong branching initialization as *pseudocost branching* (PC).

As mentioned earlier, pseudocosts play a big role in most MIP branching strategies. Because the definition of a backdoor relies on branching, it is natural to consider ways in which pseudocosts can help steer `BaMCTS` towards promising variables. Indeed, we will show a bit later how pseudocosts can serve as *global scores* for actions (variables). Those global scores are then naturally incorporated into the selection and expansion steps of MCTS. For now, we emphasize that `BaMCTS` tracks the pseudocosts resulting from each candidate evaluation. We then maintain, for each variable, a running average of its pseudocost score across all candidate evaluations that involved the variable.

Action space reduction   We would like `BaMCTS` to scale to MIPs with tens or hundreds of thousands of integer variables. However, the MCTS tree grows fast with the number of integer variables, which may hamper progress. To reduce the action space, we leverage the empirical observation that only a few of the integer variables take on fractional values in the solution of the LP relaxation of the MIP; for instance, Berthold [2014] shows that, on average across 159 instances from older MIPLIB instance libraries, 71.7% of the integer variables are integer in the LP relaxation solution. Rather than work with the full integer set I, we restrict the action space to the subset $I_{\text{frac}} \subseteq I$ of variables that are fractional in the LP relaxation. While heuristic, this restriction has some grounds in empirical MIP solving and reduces the action space dramatically.

4.2   Instantiating the four steps

Selection   We adopt a variant of the UCT selection rule, inspired by Gaudel and Sebag [2010], that adds a global action score (average pseudocosts in our case) to UCT's typical elements. Consider an MCTS search tree node (or state) S whose child node S' is being assessed; assume S' extends S with variable $i \in I_{\text{frac}}$.

We let T_S denote the number of visits to node S; $\text{EXP}(S, S')$ denotes the exploration score of S', which is large when $T_{S'}$ is much smaller than T_S; $\hat{\mu}_{S'}$ denotes the current average reward of state S'; $\mathbf{r}_{S'}$ is the vector of rewards that have been observed in the subtree rooted at S'. Our final scoring function for node selection is given by (1). With $\alpha_{\text{PC}} \in [0, 1)$, the scoring function is a convex combination of a UCT-type score for state S' and the average pseudocost score $\hat{\text{PC}}_i$ of variable i. The latter may be interpreted as a RAVE score following [Gelly and Silver, 2007]. To arrive at the final scoring function, we define the exploration score (based on the standard UCT formula), the variance score (based on the UCB1-Tuned of Auer et al. [2002]), UCT without variance, UCT with variance, and the UCT score which is one of the two preceding scores depending on the value of `use_variance`, respectively:

$$\text{EXP}(S, S') = \sqrt{\frac{\ln(T_S)}{T_{S'}}}$$

$$\text{VAR}(S') = \sqrt{\min\left\{\frac{1}{4}, \text{variance}(\mathbf{r}_{S'}) + \text{EXP}(S, S')\right\}}$$

$$\text{UCT}(S, S') = \hat{\mu}_{S'} + C \cdot \text{EXP}(S, S')$$

$$\text{UCT}_{\text{var}}(S, S') = \hat{\mu}_{S'} + C \cdot \text{EXP}(S, S') \cdot \text{VAR}(S')$$

$$\text{UCT}_{\text{score}}(S, S') = \begin{cases} \text{UCT}_{\text{var}}(S, S') & \text{if } \texttt{use\_variance}, \\ \text{UCT}(S, S') & \text{otherwise.} \end{cases}$$

The final scoring function then writes:

$$\text{SCORE}(S, S') = (1 - \alpha_{\text{PC}}) \cdot \text{UCT}_{\text{score}}(S, S') + \alpha_{\text{PC}} \cdot \hat{\text{PC}}_i \tag{1}$$

To conclude, we note that the scoring function has three hyperparameters whose effects will be analyzed experimentally in the next section:

- $\alpha_{\text{PC}} \in [0, 1)$: the weight accorded to the global pseudocost average;
- $C \in \mathbb{R}_{>0}$: the exploration weight;
- `use_variance` $\in \{\text{True}, \text{False}\}$: a boolean that determines whether UCT with or without variance is used.

Expansion   Besides the traditional *uniform random* expansion rule, we consider a deterministic *best score* expansion rule which simply expands using the available action (variable) with the largest average pseudocost score. Together, these two rules cover a wide range along the exploration-exploitation spectrum.

Simulation   We opt to stick with random simulation as described in the preceding section.

Backpropagation   We consider both the sum and max-backup rules here. While the latter is typically considered to be overly aggressive and wasteful (of reward information), it is quite suitable for our setting and allows for a form of focused local search: if a high-reward candidate has been observed in a node's subtree, a max-backup encourages more future visits to the same node. This may lead to slight modifications to the candidate that bring about improved rewards.

4.3   Algorithm engineering

Our implementation of `BaMCTS` exploits certain properties of the backdoor search problem to speed it up. One such property is that the root node of the branch-and-bound tree of each candidate evaluation is the same. Because solving the LP relaxation of the root node is typically much more time-consuming than other subproblems', we instrument CPLEX to solve the root LP once for all in advance and reuse its solution in subsequent evaluations. Rather than simulate and evaluate a single candidate at a time, we leverage the independence between the simulations to execute them and the candidate evaluation in parallel.
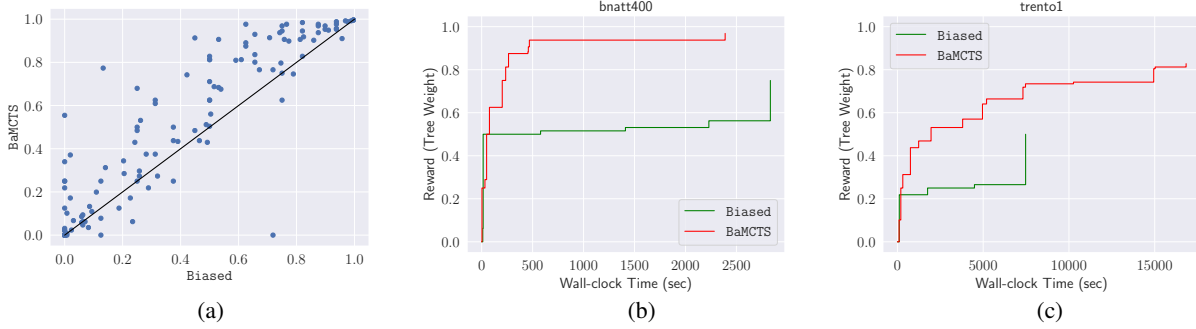
Figure 1: (a) BaMCTS outperforms Biased in finding backdoors with higher (mean) reward: most instances fall to the top of the diagonal. (b) & (c) BaMCTS attains higher rewards faster than Biased on two representative instances.
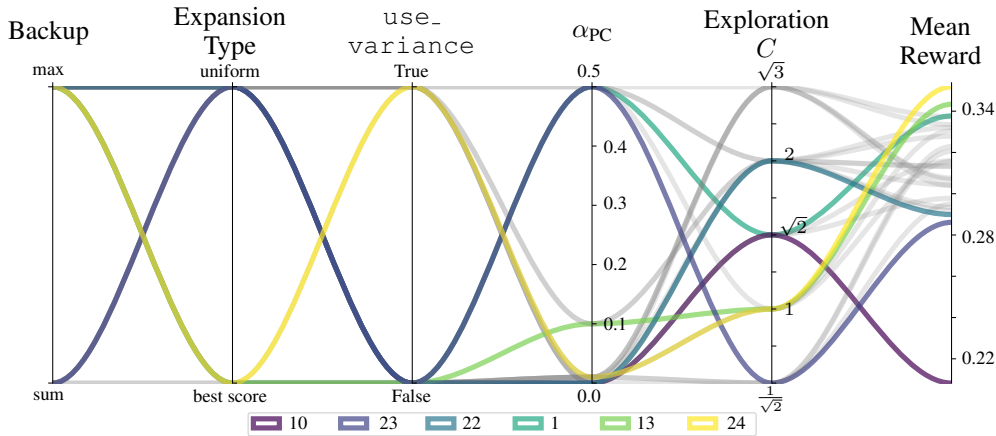

Figure 2: BaMCTS is robust wrt. hyperparameter changes. Each curve represents a hyperparameter configuration (30 total) and indicates the values taken by that config as well as the obtained mean reward after running BaMCTS with that config on the dataset for 1 hour. Notice the congregation of rewards within the [0.28, 0.35] band. This shows that for most configs the reward does not change drastically. The highlighted curves represent the best and worst performing configs. The details of the highlighted curves are shown in Table 1 and the others in the Appendix.

5   Experiments

To evaluate our method, we designed a set of experiments to answer the following questions: **1) Backdoor Extraction:** Can BaMCTS find backdoors with better tree weight values (or rewards) compared to the biased random sampling of Dilkina et al. [2009]? **2) Sensitivity Analysis:** How sensitive is BaMCTS to its hyperparameters? **3) Correlation Between Tree Weight and Branch-and-Bound Tree Size:** Is branching on higher tree weight backdoors conducive to smaller search trees or (for instances that are not solved to optimality within a time limit) a smaller optimality gap? In other words, is tree weight a legitimate reward function for BaMCTS?

5.1   Experimental Setup

Instances   We conducted our experiments on the MIPLIB2017 Benchmark set [1] [Gleixner et al., 2021] that contains 240 instances. We only considered the 164 mixed-binary instances, i.e., instances with no general integer variables, due to the ease of implementation of the tree weight for binary problems; extension to general integer variables is possible. We presolved the instances in advance to eliminate redundant variables and constraints, and also let CPLEX generate cuts at the root node; the resulting preprocessed instance was used instead of the original in all subsequent procedure, as presolving and root cuts are both standard steps in MIP solving, indepedently of any backdoor search. Some instances were excluded due to either memory issues during presolving or being solved at the root node without branching, reducing the final dataset size down to 142 instances.

---

[1]http://miplib2017.zib.de/tag_benchmark.html

	Backup	Exp. Type	use_variance	α_{PC}	C	Mean Reward
24	max	Best	T	0.01	1	0.351
13	max	Best	F	0.10	1	0.343
1	max	Best	F	0.50	$\sqrt{2}$	0.337
22	max	Uniform	F	0.00	2	0.289
23	sum	Uniform	F	0.50	$1/\sqrt{2}$	0.286
10	max	Uniform	F	0.00	$\sqrt{2}$	0.208

Table 1: Top and bottom 3 (out of 30) best performing `BaMCTS` parameter configurations in terms of mean reward (tree weight).

Baselines   We compared `BaMCTS` against two baselines: `Biased` and our implementation of `SetCover` from Fischetti and Monaci [2011]. Each run consisted of two phases: backdoor search and backdoor-guided MIP solving.

Protocol for backdoor search   For the search phase, each method was given a budget of 5 hours per instance. We recorded the sequence of improving backdoors found by each method on every instance, particularly the backdoor with the highest tree weight. For a fair comparison, we tasked all methods to find backdoors of a specific size K. Naturally, it is trivial to find large backdoors (at the extreme, the full integer set is backdoor) but their usefulness in the downstream MIP solving phase diminishes as they get larger. We report our results for backdoor size of $K = 8$; however, we experimented with other backdoor sizes in the 5-10 range and the results still carry.

Protocol for MIP solving   In the second phase, we solved the same instance for one hour using CPLEX, but instrumented the solver to prioritize branching on the backdoor variables in the order they are given. To mitigate CPLEX's randomness w.r.t. arbitrary initial conditions [Lodi and Tramontani, 2013], we solve each instance with 3 random seeds.

BaMCTS Hyperparameters   To test the sensitivity of `BaMCTS` we performed a random grid search over its hyperparameters. We selected 30 unique parameter configurations and ran `BaMCTS` for one hour per instance and per configuration, and picked the configuration with the highest mean tree weight for the 5-hour backdoor search phase introduced above. We performed the grid search over these domains:

- Backup $\in \{\text{sum}, \text{max}\}$
- Expansion Type $\in \{\text{Best Score}, \text{Uniform}\}$
- use_variance $\in \{\text{True}, \text{False}\}$
- $\alpha_{PC} \in \{0.00, 0.01, 0.10, 0.50\}$
- Exploration Parameter $(C) \in \{\frac{1}{\sqrt{2}}, 1, \sqrt{2}, 2, \sqrt{3}\}$.

Hardware   All experiments were conducted on a large CPU cluster. All MIP solving runs use a single core with an 8GB memory limit and a 1-hour time limit. Backdoor search runs for both `BaMCTS` and `Biased` use 10 cores in parallel with a variable memory limit, capped at 63GB, that is proportional to the instance's size on disk.

5.2   Results

Backdoor Extraction Performance   Figure 1 compares the performance of `BaMCTS` vs. `Biased` in terms of the quality of the backdoors they find within the 5-hour time limit per instance. `BaMCTS` clearly outperformed `Biased`, leading to discovery of higher tree weight backdoors on many more instances. This is further demonstrated in the top-right plot of Figure 4. There, we observe that `BaMCTS` with even 1 hour time budget (blue) generally outperformed `Biased` with 5 hour budget (green). `BaMCTS` with 5 hour budget (red) further widened the gap. Note that `SetCover` does not depend on tree weight as a reward and is thus not comparable with `BaMCTS` here.

Sensitivity Analysis   The box-plots of figure 4 show the per-instance sensitivity of `BaMCTS` to the set of 30 hyperparameter configurations we tested. The instances are sorted by median reward from low (hard instances) to high (easy instances). Even though we observe quite high sensitivity for some instances, for the majority of them most configurations obtain similar reward. Figure 2 shows all 30 hyperparameter configurations for `BaMCTS` as well as the average reward that each configuration achieved over the entire dataset. Most of the configurations resulted in average reward in the range of [0.28, 0.35], again indicating that at least on average there is not a huge variability in terms of the reward, and `BaMCTS` can be used out of the box with default hyperparameter values and achieve a reasonable

Table 2: Summary of MIP solving results. "seed 1 (47, 56, 4)" means that for this seed, 47 instances were solved by both CPX−def and CPX−BaMCTS, 56 were not solved by either within 1 hour, and 4 were solved by only one of the two. The better method for each metric is in bold.

	seed 1 (47, 56, 4)		seed 2 (45, 61, 7)		seed 3 (46, 60, 6)	
	CPX−def	CPX−BaMCTS	CPX−def	CPX−BaMCTS	CPX−def	CPX−BaMCTS
# of nodes (solved by both)	6902.74	**6097.70**	6173.84	**5030.13**	7744.98	**7001.99**
total time (solved by both)	179.53	**175.97**	156.61	**138.34**	169.41	**156.83**
optimality gap (not solved by either)	23/56	**33/56**	24/61	**37/61**	20/60	**40/60**
# of instances (solved by one method)	**3**	1	**4**	3	3	3

performance. Note that MIPLIB2017 is quite diverse and it is likely that testing on a more homogeneous dataset would result in even lower parameter sensitivity. Appendix Table 3 lists the reward values for all tested configurations.

Tree Weight and Branch-and-Bound Tree Size   Table 2 summarizes the MIP solving results. CPX−def is CPLEX 12.10.0 with traditional branch-and-bound (i.e., with CPLEX's "dynamic search" turned off); results with "dynamic search" are similar. CPX−BaMCTS is CPLEX 12.10.0 with branching priorities defined by the best backdoor found by the best hyperparameter configuration of BaMCTS (**24** in Table 1) for that instance. Of the initial 142 instances, we restrict the MIP solving here to 115 instances for which BaMCTS found at least one backdoor candidate with non-zero tree weight. Within each random seed, the instances are partitioned into three sets: (i) instances that are solved by both methods, for which we compare the shifted geometric means of the number of nodes and total time (lower is better; see Appendix A), with shifts 100 and 10 respectively; (ii) instances that are not solved by either method, for which the optimality gap at termination after 1 hour tells us which method made more progress; we count the number of instances for which each a method achieved a smaller gap (higher is better); and (iii) instances that are solved by exactly one of the two methods, for which we simply count the number of such instances for each method (higher is better).

Table 2 shows that CPX−BaMCTS wins on the first three metrics consistently across the three seeds, with an average reduction of 700 to 1100 in the number of nodes. The last metric (and row of the table) records two wins for CPX−def and one tie, but it is over 4 to 7 instances out of more than 110. We've also used the two-sided Wilcoxon signed-rank test, as described in [Hendel, 2014], to compare the distribution of values for these metrics, and have found the corresponding p-values to be typically small, implying that the observed values for these metrics come from distributions with different medians for each of CPX−def and CPX−BaMCTS; this agrees with the conclusions from Table 2.

The fact that branching with the best backdoor found by BaMCTS typically leads to better MIP solving (smaller number of nodes, smaller optimality gaps, etc.) confirms that tree weight is an informative reward signal for finding high-quality backdoor candidates.

Lastly, we note that we have attempted to compare against the SetCover method of Fischetti and Monaci [2011] on the MIP solving metrics. However, SetCover fails to return a size-8 backdoor for half of the 142 instances, and requires a larger number of variables to cover fractional vertices. This makes it difficult to compare BaMCTS to SetCover. To address this limitation of SetCover, modifications to the algorithm may be necessary but are out of scope for this paper.

Impact on total time   While Table 2 shows that branching with the backdoors found by our method result in only small reductions in total time, we note that: (1) this reduction is consistent across the three seeds and is in the 2-10% range, which is non-trivial for standard MIPLIB2017 Benchmark instances; and (2) the experiment aims primarily at showing that branching with the backdoors translates into a smaller *number of nodes*, a metric for which improvements are larger. Because CPLEX is not open-source, we are unable to control its behavior beyond its programming interface (i.e., callbacks), but believe that further total time speedups may be achieved with tighter integration with the internal solver code.

Tree weight vs. Number of variables Give the wide range of tree weight values that were observed as a result of backdoor search, we were interested in whether only instances with a small action space, i.e., a small number of integer variables that are fractional in the solution of the MIP LP relaxation at the root node of branch-and-bound tree, were associated with tree weights (rewards) that are close to 1. Figure 3 shows that is not the case at all: the tree weight values are spread across the wide range of values on the horizontal axis, and so our results are not biased towards "easy" small instances.


Figure 3: Scatter plot of the number of fractional variables in the MIP root LP relaxation solution vs. the best reward (i.e., tree weight) found by our method.

6 Conclusion

In this paper we proposed `BaMCTS`, a Monte Carlo Tree Search framework for finding strong backdoors in MIPs and demonstrated through our experiments its merits relative to earlier backdoor search algorithms in MIP literature. We hope that this paper would motivate further research in this direction and position MCTS as a viable approach in finding MIP backdoors. One particularly interesting future direction is to apply this technique on more homogeneous families of instances instead of MIPLIB to discover inherent structures for those problem families by studying their backdoors. We are planning to release our code to streamline these efforts.

Figure 4: Per-instance reward distribution. *main*: Box-plots showing the range and median of rewards for each instance over the set of 30 hyperparameter configs. Instances are sorted based on median reward from hard (red) to easy (blue). *top-right*: Comparison between `Biased` and two `BaMCTS` runs, with 1 & 5 hour budgets.

References

Bistra Dilkina, Carla P Gomes, Yuri Malitsky, Ashish Sabharwal, and Meinolf Sellmann. Backdoors to combinatorial optimization: Feasibility and optimality. In *International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, pages 56–70. Springer, 2009.

Tobias Achterberg, Thorsten Koch, and Alexander Martin. Miplib 2003. *Operations Research Letters*, 34(4):361–372, 2006.

Matteo Fischetti and Michele Monaci. Backdoor branching. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 183–191. Springer, 2011.

Matteo Fischetti and Michele Monaci. Exploiting erraticism in search. *Operations Research*, 62(1):114–122, 2014.

Matteo Fischetti. Branchstorming (brainstorming about tree search). URL: http://www.dei.unipd.it/~fisch/papers/slides/2014%20ISCO%20%5bplenary%20Fischetti%20BRANCHstorming%5d.pdf, March 2014. Plenary lecture at ISCO 2014.

Ambros Gleixner, Gregor Hendel, Gerald Gamrath, Tobias Achterberg, Michael Bastubbe, Timo Berthold, Philipp Christophel, Kati Jarck, Thorsten Koch, Jeff Linderoth, et al. Miplib 2017: data-driven compilation of the 6th mixed-integer programming library. *Mathematical Programming Computation*, pages 1–48, 2021.

Ryan Williams, Carla P Gomes, and Bart Selman. Backdoors to typical case complexity. In *IJCAI*, volume 3, pages 1173–1178, 2003.

Lionel Paris, Richard Ostrowski, Pierre Siegel, and Lakhdar Sais. Computing horn strong backdoor sets thanks to local search. In *2006 18th IEEE International Conference on Tools with Artificial Intelligence (ICTAI'06)*, pages 139–143. IEEE, 2006.

Stephan Kottler, Michael Kaufmann, and Carsten Sinz. Computation of renameable horn backdoors. In *International Conference on Theory and Applications of Satisfiability Testing*, pages 154–160. Springer, 2008.

Zijie Li and Peter Van Beek. Finding small backdoors in sat instances. In *Canadian Conference on Artificial Intelligence*, pages 269–280. Springer, 2011.

David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550 (7676):354–359, 2017.

Ashish Sabharwal, Horst Samulowitz, and Chandra Reddy. Guiding combinatorial optimization with uct. In *International conference on integration of artificial intelligence (AI) and operations research (OR) techniques in constraint programming*, pages 356–361. Springer, 2012a.

Baba Satomi, Yongjoon Joe, Atsushi Iwasaki, and Makoto Yokoo. Real-time solving of quantified csps based on monte-carlo game tree search. In *Twenty-Second International Joint Conference on Artificial Intelligence*, 2011.

Dimitris Bertsimas, J Daniel Griffith, Vishal Gupta, Mykel J Kochenderfer, Velibor V Mišić, and Robert Moss. A comparison of monte carlo tree search and mathematical optimization for large scale dynamic resource allocation. *arXiv preprint arXiv:1405.5498*, 2014.

Arpad Rimmel, Fabien Teytaud, and Tristan Cazenave. Optimization of the nested monte-carlo algorithm on the traveling salesman problem with time windows. In *European Conference on the Applications of Evolutionary Computation*, pages 501–510. Springer, 2011.

Manuel Loth, Michele Sebag, Youssef Hamadi, and Marc Schoenauer. Bandit-based search for constraint programming. In *International Conference on Principles and Practice of Constraint Programming*, pages 464–480. Springer, 2013.

Kenshin Abe, Zijian Xu, Issei Sato, and Masashi Sugiyama. Solving np-hard problems on graphs with extended alphago zero. *arXiv preprint arXiv:1905.11623*, 2019.

Ruiyang Xu, Prashank Kadam, and Karl Lieberherr. First-order problem solving through neural mcts based reinforcement learning. *arXiv preprint arXiv:2101.04167*, 2021.

William Cook. Markowitz and manne+ eastman+ land and doig= branch and bound. *Optimization Stories*, pages 227–238, 2012.

George L Nemhauser and Laurence A Wolsey. *Integer and Combinatorial Optimization*. John Wiley & Sons, 1988.

Tobias Achterberg, Thorsten Koch, and Alexander Martin. Branching rules revisited. *Operations Research Letters*, 33 (1):42–54, 2005.

Tobias Achterberg and Timo Berthold. Hybrid branching. In *CPAIOR*, pages 309–311. Springer, 2009.

Gregor Hendel. Enhancing mip branching decisions by using the sample variance of pseudo costs. In *Integration of AI and OR Techniques in Constraint Programming*, volume 9075, pages 199 – 214, 2015. doi: 10.1007/978-3-319-18008-3_14. in press.

Cameron B Browne, Edward Powley, Daniel Whitehouse, Simon M Lucas, Peter I Cowling, Philipp Rohlfshagen, Stephen Tavener, Diego Perez, Spyridon Samothrakis, and Simon Colton. A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games*, 4(1):1–43, 2012.

Levente Kocsis and Csaba Szepesvári. Bandit based monte-carlo planning. In *ECML*. 2006.

Rémi Coulom. Computing "elo ratings" of move patterns in the game of go. *ICGA journal*, 30(4):198–208, 2007.

Adrien Couëtoux, Jean-Baptiste Hoock, Nataliya Sokolovska, Olivier Teytaud, and Nicolas Bonnard. Continuous upper confidence trees. In *International Conference on Learning and Intelligent Optimization*, pages 433–445. Springer, 2011.

Ashish Sabharwal, Horst Samulowitz, and Chandra Reddy. Guiding combinatorial optimization with uct. In *Integration of AI and OR Techniques in Contraint Programming for Combinatorial Optimzation Problems*, pages 356–361. Springer, 2012b.

Philip Kilby, John Slaney, Sylvie Thiébaux, Toby Walsh, et al. Estimating search tree size. In *Proc. of the 21st National Conf. of Artificial Intelligence, AAAI, Menlo Park*, 2006.

Gregor Hendel, Daniel Anderson, Pierre Le Bodic, and Marc Pfetsch. Estimating the size of branch-and-bound trees. 2020.

Jeff T Linderoth and Martin WP Savelsbergh. A computational study of search strategies for mixed integer programming. *INFORMS Journal on Computing*, 11(2):173–187, 1999.

Timo Berthold. Rens. *Mathematical Programming Computation*, 6(1):33–54, 2014.

Romaric Gaudel and Michele Sebag. Feature selection as a one-player game. In *International Conference on Machine Learning*, pages 359–366, 2010.

Sylvain Gelly and David Silver. Combining online and offline knowledge in uct. In *Proceedings of the 24th international conference on Machine learning*, pages 273–280, 2007.

Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2):235–256, 2002.

Andrea Lodi and Andrea Tramontani. Performance variability in mixed-integer programming. *Tutorials in Operations Research: Theory Driven by Influential Applications*, pages 1–12, 2013.

Gregor Hendel. Empirical analysis of solving phases in mixed integer programming. Master's thesis, Technische Universität Berlin, 2014.

A Further definitions

A.1 Shifted Geometric Means in Table 2

Compared to arithmetic means, geometric means are more robust to measurements that are extremely small or large. That is why they are the standard method for aggregating both the number of (branch-and-bound) nodes and total running time in MIP solving. The *shifts* are used to eliminate large *relative variations* between two methods, even when in absolute terms these variations are likely insignificant, e.g., for the same instance, two solvers A and B record running times of 0.8 and 1 second which suggests a substantial 20% improvement for A over B, whereas wall-clock times can exhibit similar fluctuations for reasons not related to the actual solvers themselves. We refer to page 33 of [Hendel, 2015] for further discussion.

The shifted geometric mean with shift τ for a set of measurements $(X_i)_{i=1}^q$ is:

$$\bar{X}^\tau = \left(\prod_{i=1}^q (X_i + \tau) \right)^{\frac{1}{q}} - \tau.$$

We use the standard shifts of $\tau = 10$ seconds for total time and $\tau = 100$ nodes for total nodes.

B Additional Experimental Results

B.1 Full `BaMCTS` Hyperparameter Tuning Results

	Backup	Exp. Type	use_variance	α_{PC}	C	Mean Reward
10	max	Uniform	False	0.00	$\sqrt{2}$	0.208
23	sum	Uniform	False	0.50	$1/\sqrt{2}$	0.286
22	max	Uniform	False	0.00	2	0.289
20	sum	Uniform	True	0.10	2	0.291
11	max	Uniform	True	0.00	2	0.293
19	max	Uniform	False	0.01	$1/\sqrt{2}$	0.295
21	sum	Uniform	False	0.00	$\sqrt{2}$	0.297
16	sum	Uniform	False	0.50	$\sqrt{3}$	0.304
12	max	Uniform	True	0.50	2	0.304
7	max	Uniform	True	0.00	$\sqrt{3}$	0.306
3	max	Uniform	False	0.01	$\sqrt{3}$	0.307
9	max	Uniform	True	0.10	2	0.313
0	sum	Uniform	False	0.01	2	0.313
17	sum	Uniform	False	0.50	1	0.315
6	sum	Uniform	True	0.00	$\sqrt{2}$	0.315
26	max	Uniform	True	0.01	$1/\sqrt{2}$	0.320
2	max	Uniform	True	0.00	$1/\sqrt{2}$	0.322
4	sum	Best	False	0.00	2	0.327
25	max	Best	False	0.01	1	0.330
15	max	Best	True	0.50	2	0.332
5	max	Best	False	0.01	$\sqrt{3}$	0.333
8	sum	Best	False	0.01	$\sqrt{3}$	0.337
1	max	Best	False	0.50	$\sqrt{2}$	0.337
13	max	Best	False	0.10	1	0.343
24	max	Best	True	0.01	1	0.351

Table 3: All `BaMCTS` parameter configurations sorted by mean reward (tree weight); higher is better.