

Improved Sample Complexity Bounds for Branch-And-Cut

Maria-Florina Balcan ✉

Computer Science and Machine Learning Departments,
Carnegie Mellon University, Pittsburgh, PA, USA

Siddharth Prasad ✉

Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, USA

Tuomas Sandholm ✉

Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, USA
Optimized Markets, Inc., Pittsburgh, PA, USA
Strategic Machine, Inc., Pittsburgh, PA, USA
Strategy Robot, Inc., Pittsburgh, PA, USA

Ellen Vitercik ✉

Department of Electrical Engineering and Computer Sciences,
University of California Berkeley, CA, USA

Abstract

The branch-and-cut algorithm for integer programming has a wide variety of tunable parameters that have a huge impact on its performance, but which are challenging to tune by hand. An increasingly popular approach is to use machine learning to configure these parameters based on a training set of integer programs from the application domain. We bound how large the training set should be to ensure that for any configuration, its average performance over the training set is close to its expected future performance. Our guarantees apply to parameters that control the most important aspects of branch-and-cut: node selection, branching constraint selection, and cut selection, and are sharper and more general than those from prior research.

2012 ACM Subject Classification Theory of computation → Integer programming; Theory of computation → Sample complexity and generalization bounds

Keywords and phrases Automated algorithm configuration, integer programming, machine learning theory, tree search, branch-and-bound, branch-and-cut, cutting planes, sample complexity, generalization guarantees, data-driven algorithm design

Digital Object Identifier 10.4230/LIPIcs.CP.2022.3

Funding This material is based on work supported by the National Science Foundation under grants CCF-1733556, CCF-1910321, IIS-1901403, and SES-1919453, the ARO under award W911NF2010081, the Defense Advanced Research Projects Agency under cooperative agreement HR00112020003, a Simons Investigator Award, an AWS Machine Learning Research Award, an Amazon Research Award, a Bloomberg Research Grant, and a Microsoft Research Faculty Fellowship.

1 Introduction

Branch-and-cut (B&C) is a powerful algorithmic paradigm that is the backbone of all modern integer programming (IP) solvers. The main components of B&C can be tuned and tweaked in a myriad of ways. The fastest commercial IP solvers like CPLEX and Gurobi employ an array of heuristics to make decisions at every stage of B&C to reduce the solving time as much as possible, and give the user freedom to tune the multitude of parameters influencing the search through the space of feasible solutions. However, tuning the parameters that control B&C in a principled way is an inexact science with little to no formal mathematical guidelines. A rapidly growing line of work studies machine-learning approaches to speeding



© Maria-Florina Balcan, Siddharth Prasad, Tuomas Sandholm, and Ellen Vitercik;
licensed under Creative Commons License CC-BY 4.0

28th International Conference on Principles and Practice of Constraint Programming (CP 2022).

Editor: Christine Solnon; Article No. 3; pp. 3:1–3:19



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

up the various aspects of B&C – in particular investigating whether high-performing B&C parameter configurations can be learned from a *training set* of typical IPs from the particular application at hand [2, 22, 37, 42, 25, 32, 27, 43, 29]. Complementing the substantial number of experimental approaches using machine learning for B&C, a recent generalization theory has developed in parallel that aims to provide a rigorous theoretical foundation for how well any B&C configuration learned from training IP data will perform on new unseen IPs [7, 9]. In particular, this line of theoretical research provides *sample complexity guarantees* that bound how large the training set should be to ensure that *no matter how the parameters are configured* (i.e., using any approach from prior research), the average performance of branch-and-cut over the training set is close to its expected future performance. Sample complexity bounds are important because with too small a training set, learning is impossible: a configuration may have strong average performance over the training set but terrible expected performance on future IPs. If the training set is too small, then no matter how the parameters are tuned, the resulting configuration will not have reliably better performance than any default configuration. State-of-the-art parameter tuning methods have historically come without any provable guarantees, and our results fill in that gap for a wide array of tunable B&C parameters. In this paper, we expand and improve upon the existing theory to develop a wider and sharper handle on the learnability of the key components of B&C.

1.1 Summary of main contributions

Our main contribution is a formalization of a general model of tree search, presented in Section 2.1, that allows us to improve and generalize prior results on the sample complexity of tuning B&C. In this model, the algorithm repeatedly chooses a leaf node of the search tree, performs a series of actions (for example, a cutting plane to apply and a constraint to branch on), and adds children to that leaf in the search tree. The algorithm will also fathom nodes when applicable. The node and action selection are governed by *scoring rules*, which assign a real-valued score to each node and possible action. For example, a node-selection scoring rule might equal the objective value of the node’s LP relaxation. We focus on general tree search with *path-wise* scoring rules. At a high level, a score of a node or action is path-wise if its value only depends on information contained along the path between the root and that node, as is often the case in B&C. Many commonly used scoring rules are path-wise including the efficacy [4], objective parallelism [1], directed cutoff distance [15], and integral support [41] scoring rules, all used for cut selection by the leading open-source solver SCIP [15]; the best-bound scoring rule for node selection; and the linear, product, and most-fractional scoring rules for variable selection using strong branching [1]. In Section 4, we show how this general model of tree search captures a wide array of B&C components, including node selection, general branching constraint selection, and cutting plane selection, simultaneously. We also provide experimental evidence that, in the case of cutting plane selection, the data-dependent tuning suggested by our model can lead to dramatic reductions in the number of nodes expanded by B&C.

In Section 3, we prove our main structural result: for any IP, the tree search parameter space can be partitioned into a finite number of regions such that in any one region, the resulting search tree is fixed. This is in spite of the fact that the B&C search tree can be an extremely unstable function of its parameters, with minuscule changes leading to exponentially better or worse performance [7, 9]. By analyzing the complexity of this partition, we prove our sample complexity bound. In particular, we relate the complexity of the partition to the *pseudo-dimension* of the set of functions that measure the performance of B&C as a function of the input IP. Pseudo-dimension (defined in Section 3) is a combinatorial notion

from machine learning theory that measures the *intrinsic complexity* of a set of functions. At a high level, it measures how well a set of functions are able match complex patterns. Classic results from learning theory then allow us to translate our pseudo-dimension bound into a sample complexity guarantee [3], capturing the intuition that the more complex patterns one can fit (i.e., the larger the pseudo-dimension is), the more samples needed to generalize. The sample complexity bound grows linearly with the pseudo-dimension, so ideally, the pseudo-dimension will be polynomial in the size of the problem.

We show that the pseudo-dimension is only polynomial in the depth of the tree (which is, for example, at most the number of variables in the case of binary integer programming). By contrast, we might naïvely expect the pseudo-dimension to grow linearly with the number of arithmetic operations required to compute the B&C tree (as in Theorem 8.4 by Anthony and Bartlett [3]), which is exponential in the depth of the tree. In fact, our bound is exponentially smaller than the pseudo-dimension bound of prior research by Balcan et al. [9], which grows linearly with the total number of nodes in the tree. Their results apply to any type of scoring rule, path-wise or otherwise. By taking advantage of the path-wise structure, we are able to reason inductively over the depth of the tree, leading to our exponentially improved bound. Our results recover those of Balcan et al. [7], who only studied path-wise scoring rules for single-variable selection for branching. In contrast, we are able to handle many more of the critical components of tree search: node selection, general branching constraint selection, and cutting plane selection.

1.2 Additional related research

A growing body of research has studied how machine learning can be used to speed up the time it takes to solve integer programs, primarily from an empirical perspective, whereas we study this problem from a theoretical perspective. This line of research has included general parameter tuning procedures [25, 27, 24, 37], which are not restricted to any one aspect of B&C. Researchers have also honed in on specific aspects of tree search and worked towards improving those using machine learning. These include variable selection [29, 2, 13, 7, 16, 19], general branching constraint selection [44], cut selection [37, 39, 23, 9], node selection [36, 21], and heuristic scheduling [30, 10]. Machine learning approaches to large neighborhood search have also been used to speed up solver runtimes [38].

This paper contributes to a line of research that provides sample complexity guarantees for algorithm configuration, often by using structure exhibited by the algorithm’s performance as a function of its parameters [20, 8, 7, 6, 9, 5]. This line of research has studied algorithms for clustering [8], computational biology [6], and integer programming [7, 9], among other computational problems. The main contribution of this paper is to provide a sharp yet general analysis of the performance of tree search as a function of its parameters.

A related line of research provides algorithm configuration procedures with provable guarantees that are agnostic to the specific algorithm that is being configured [31, 40] and are particularly well-suited for algorithms with a finite number of possible configurations (though they can be applied to algorithms with infinite parameter spaces by randomly sampling a finite set of configurations).

2 Main tree search model

In this section we present our general tree search model and situate it within the framework of sample complexity. Balcan et al. [9] studied the sample complexity of a much more general formulation of a tunable search algorithm without any inherent tree structure. Our formulation explicitly builds a tree.

Algorithm 1 Tree search.

Input: Root node Q , depth limit Δ

```

1: Initialize  $\mathcal{T} = Q$ .
2: while  $\mathcal{T}$  contains an unfathomed leaf do
3:   Select a leaf  $Q$  of  $\mathcal{T}$  that maximizes  $\text{nscore}(\mathcal{T}, Q)$ .
4:   if  $\text{depth}(Q) = \Delta$  or  $\text{fathom}(\mathcal{T}, Q, \text{None})$  then
5:     Fathom  $Q$ .
6:   else
7:     Select an action  $A \in \text{actions}(\mathcal{T}, Q)$  that maximizes  $\text{ascore}(\mathcal{T}, Q, A)$ .
8:     if  $\text{fathom}(\mathcal{T}, Q, A)$  then
9:       Fathom  $Q$ .
10:    else if  $\text{children}(\mathcal{T}, Q, A) = \emptyset$  then
11:      Fathom  $Q$ .
12:    else
13:      Add all nodes in  $\text{children}(\mathcal{T}, Q, A)$  to  $\mathcal{T}$  as children of  $Q$ .

```

2.1 General model of tree search

Tree search starts with a root node. In each round of tree search, a leaf node Q is selected. At this node, one of three things may occur: (1) Q is fathomed, meaning it is never visited again, (2) some action is taken at Q , and then it is fathomed, or (3) some action is taken at Q , and then some number of children nodes of Q are added to the tree. (For example, an action might represent a decision about which variable to branch on.) This process repeats until the tree has no unfathomed leaves. More formally, there are functions **actions**, **children**, and **fathom** prescribing how the search proceeds. Given a partial tree \mathcal{T} and a leaf Q of \mathcal{T} , $\text{actions}(\mathcal{T}, Q)$ outputs a set of actions available at Q . Given a partial tree \mathcal{T} , a leaf Q of \mathcal{T} , and an action $A \in \text{actions}(\mathcal{T}, Q)$, $\text{fathom}(\mathcal{T}, Q, A) \in \{\text{true}, \text{false}\}$ is a Boolean function used to determine when to fathom a leaf Q of \mathcal{T} given that action $A \in \text{actions}(\mathcal{T}, Q) \cup \{\text{None}\}$ was taken at Q , and $\text{children}(\mathcal{T}, Q, A)$ outputs a (potentially empty) list of nodes representing the children of Q to be added to the search tree given that action A was taken at Q . Finally, $\text{nscore}(\mathcal{T}, Q)$ is a node-selection score that outputs a real-valued score for each leaf of \mathcal{T} , and $\text{ascore}(\mathcal{T}, Q, A)$ is an action-selection score that outputs a real-valued score for each action $A \in \text{actions}(\mathcal{T}, Q)$. These scores are heuristics that are meant to indicate the quality of exploring a node or performing an action.

Many aspects of B&C are governed by scoring rules [1]. For example, commonly used scoring rules for cutting plane selection include *efficacy* [4], which is the perpendicular distance from the current LP solution to the cutting plane; *parallelism* [1], which measures the angle between the objective and the normal vector to the cutting plane; and *directed cutoff* [15], which is the distance from the current LP solution to the cutting plane along the direction of the line segment connecting the LP solution to the current best incumbent integer solution. For node selection, under the commonly used best-first node selection policy, $\text{nscore}(\mathcal{T}, Q)$ equals the objective value of the LP relaxation of the IP represented by the node Q . Finally, for variable selection, popular scoring rules include a maximum change in LP objective value after branching on the variable (where the maximum is taken over the two resulting children), the minimum change in the LP objective value, linear combinations of these two values, and the product of these two values [1]. Algorithm 1 is a formal description of tree search using these functions.

The key condition that enables us to derive stronger sample complexity bounds compared to prior research is the notion of a *path-wise* function, which was also used in prior research but only in the context of variable selection [7].

► **Definition 1** (Path-wise functions). *A function f on tree-leaf pairs is path-wise if for all \mathcal{T} and $Q \in \mathcal{T}$, $f(\mathcal{T}, Q) = f(\mathcal{T}_Q, Q)$, where \mathcal{T}_Q is the path from the root of \mathcal{T} to Q . A function g on tree-leaf-action triples is path-wise if for all A , $f_A(\mathcal{T}, Q) := g(\mathcal{T}, Q, A)$ is path-wise.*

We assume that **actions**, **ascore**, **nscore** and **children** are path-wise, though **fathom** is not necessarily path-wise.

Many commonly-used scoring rules are path-wise. For example, scoring rules are often functions of the LP relaxation of the IP represented by a given node, and these scoring rules are path-wise. Specific examples include the efficacy, objective parallelism, directed cutoff distance, and integral support scoring rules used for cut selection; the best-bound scoring rule for node selection; and the linear, product, and most-fractional scoring rules for variable selection using strong branching. A point of clarification: the pathwise assumption is with respect to the numerical scores assigned to actions/nodes. The actual act of, for example, node selection, can depend on the entire tree. For example, consider the best-bound node selection rule in branch-and-cut, which chooses the node with the best LP estimate. Here, the scoring rule, which is the LP objective value itself, is pathwise, but ultimately the node that is selected depends on the LP bounds at every unexplored node of the tree. This is fine for our analysis. Similarly, the decision to fathom a node based on LP bounds is a decision that depends on the entire tree built so far, which is also captured by our analysis.

No one scoring rule is optimal across all application domains, and prior research on variable selection has shown that it can be advantageous to adapt the scoring rule to the application domain at hand [7]. To this end, Algorithm 1 can be tuned by two parameters $\mu \in [0, 1]$ and $\lambda \in [0, 1]$ that control action selection and node selection, respectively. Given two fixed path-wise action-selection scores **ascore**₁ and **ascore**₂, we define a new score by

$$\text{ascore}_\mu(\mathcal{T}, Q) = \mu \cdot \text{ascore}_1(\mathcal{T}, Q) + (1 - \mu) \cdot \text{ascore}_2(\mathcal{T}, Q).$$

Similarly, given two path-wise node-selection scores **nscore**₁ and **nscore**₂, we define

$$\text{nscore}_\lambda(\mathcal{T}, Q, A) = \lambda \cdot \text{nscore}_1(\mathcal{T}, Q, A) + (1 - \lambda) \cdot \text{nscore}_2(\mathcal{T}, Q, A).$$

Then, if **nscore** _{λ} and **ascore** _{μ} are used as the scores in Algorithm 1, we can view the behavior of tree search as a function of μ and λ . The choice to use a convex combination of scores is not new: prior research has shown that this idea can lead to dramatic improvements in the case of single-variable branching [7]. Furthermore, the leading open source solver SCIP uses a hard-coded weighted sum of scoring rules to select cutting planes. More broadly, interpolating between two scores is a commonly-studied modeling choice in other machine learning topics such as clustering [8].

Finally, we assume there exists $b, k \in \mathbb{N}$ such that $|\text{actions}(\mathcal{T}, Q)| \leq b$ for any $Q \in \mathcal{T}$, and $|\text{children}(\mathcal{T}, Q, A)| \leq k$ for all Q, A .

2.2 Problem formulation

We now define the notion of a *sample complexity bound* more formally. Let \mathcal{Q} denote the domain of possible input root nodes Q to Algorithm 1 (for example, the set of all IPs with n variables and m constraints). As is common in prior research on algorithm configuration [22, 37, 42, 25, 32, 27, 43], we assume there is some unknown distribution \mathcal{D}

over \mathcal{Q} . In the IP setting, \mathcal{D} could represent, for example, typical scheduling IP instances solved by an airline company. The *sample complexity* of a class of real valued functions $\mathcal{F} = \{f : \mathcal{Q} \rightarrow \mathbb{R}\}$ is the minimum number of independent samples required from \mathcal{D} so that with high probability over the samples, the empirical value of f on the samples is a good approximation of the expected value of f over \mathcal{D} , uniformly over all $f \in \mathcal{F}$. Formally, given an error parameter ε and confidence parameter δ , the sample complexity $N_{\mathcal{F}}(\varepsilon, \delta)$ is the minimum $N_0 \in \mathbb{N}$ such that for any $N \geq N_0$,

$$\Pr_{Q_1, \dots, Q_N \sim \mathcal{D}} \left(\sup_{f \in \mathcal{F}} \left| \frac{1}{N} \sum_{i=1}^N f(Q_i) - \mathbb{E}_{Q \sim \mathcal{D}}[f(Q)] \right| \leq \varepsilon \right) \geq 1 - \delta$$

for all distributions \mathcal{D} supported on \mathcal{Q} . Equivalently, our results bound the error $\varepsilon_{\mathcal{F}}(N, \delta)$ between the empirical value of any $f \in \mathcal{F}$ and its true expected value in terms of the number of training samples N and the confidence parameter δ . $N_{\mathcal{F}}(\varepsilon, \delta)$ is the number of samples required to achieve a prescribed error bound ε , while $\varepsilon_{\mathcal{F}}(N, \delta)$ provides an error bound for any number N of samples at hand. We provide bounds on $N_{\mathcal{F}}(\varepsilon, \delta)$ and $\varepsilon_{\mathcal{F}}(N, \delta)$ in terms of a common learning-theoretic measure of intrinsic complexity of \mathcal{F} called *pseudo-dimension*, which is detailed in Section 3.

In the context of Algorithm 1, we study families of *tree-constant* cost functions. A cost function $\text{cost} : \mathcal{Q} \rightarrow \mathbb{R}$ is tree constant if $\text{cost}(Q)$ only depends on the tree built by Algorithm 1 on input Q (an example is tree size). Let $\text{cost}_{\mu, \lambda}(Q)$ denote this cost when Algorithm 1 is run using the scores $\text{ascore}_{\mu} = \mu \cdot \text{ascore}_1 + (1 - \mu) \cdot \text{ascore}_2$ and $\text{nscore}_{\lambda} = \lambda \cdot \text{nscore}_1 + (1 - \lambda) \cdot \text{nscore}_2$. We study the sample complexity of $\mathcal{F} = \{\text{cost}_{\mu, \lambda} : \mu, \lambda \in [0, 1]\}$. We emphasize that we primarily interpret tree-constant functions as proxies for run-time. In the context of integer programming, tree size is one such measure. A strength of these guarantees is that they apply no matter how the parameters are tuned: optimally or suboptimally, manually or automatically. For *any* configuration, these guarantees bound the difference between average performance over the training set and expected future performance on unseen IPs.

3 Generalization guarantees for tree search

In order to derive our sample complexity guarantees, we first prove a key structural property: the behavior of Algorithm 1 is piecewise constant as a function of the node-selection score parameter λ and the action-selection score parameter μ . We give a high-level outline of our approach. We first assume that the conditional checks $\text{fathom}(\mathcal{T}, Q, \cdot) = \text{true}$ (lines 4 and 8) are suppressed. Let \mathcal{A}' denote Algorithm 1 without these checks (so \mathcal{A}' fathoms a node if and only if the depth limit is reached or if the node has no children). The behavior of \mathcal{A}' as a function of μ and λ can be shown to be piecewise constant using the same argument as in Claim 3.4 of Balcan et al. [7]. Given this, our first main technical contribution (Lemma 2) is a generalization of Claim 3.5 of Balcan et al. [7] that relates the behavior of \mathcal{A}' to Algorithm 1. The argument in Balcan et al. [7] is specific to branching, but we are able to prove our result in a much more general setting. Our second main technical contribution (Lemma 4) is to establish piecewise structure when the node-selection score is controlled by $\lambda \in [0, 1]$. The main reason for this auxiliary step of analyzing \mathcal{A}' is due to the fact that fathom is *not* necessarily a path-wise function, and can depend on the state of the entire tree.

► **Lemma 2.** *Fix $\mu \in [0, 1]$. Let \mathcal{T} and \mathcal{T}' be the trees built by Algorithm 1 and \mathcal{A}' , respectively, using the action-selection score $\mu \cdot \text{ascore}_1 + (1 - \mu) \cdot \text{ascore}_2$. Let Q be any node in \mathcal{T} , and let \mathcal{T}_Q be the path from the root of \mathcal{T} to Q . Then, \mathcal{T}_Q is a rooted subtree of \mathcal{T}' , no matter what node selection policy is used.*

Proof. Let t denote the length of the path \mathcal{T}_Q . Let \mathcal{T}_Q be comprised of the sequence of nodes (Q_1, \dots, Q_t) such that Q_1 is the root of \mathcal{T} , $Q_t = Q$, and for each τ , $Q_{\tau+1} \in \text{children}(\mathcal{T}_{Q_\tau}, Q_\tau, A_\tau)$ where $A_\tau \in \text{actions}(\mathcal{T}_{Q_\tau}, Q_\tau)$ is the action selected by Algorithm 1 at node Q_τ . We show that (Q_1, \dots, Q_t) is a rooted path in \mathcal{T}' as well.

Suppose for the sake of contradiction that this is not the case. Let $\tau \in \{2, \dots, t\}$ be the minimal index such that $(Q_1, \dots, Q_{\tau-1})$ is a rooted path in \mathcal{T}' , but there is no edge in \mathcal{T}' from $Q_{\tau-1}$ to node Q_τ . There are two possible cases:

Case 1. $Q_{\tau-1}$ was fathomed by \mathcal{A}' . This case is trivially not possible since whenever \mathcal{A}' fathoms a node, so does Algorithm 1 (recall \mathcal{A}' was defined by suppressing fathoming conditions of Algorithm 1).

Case 2. $Q_\tau \notin \text{children}(\mathcal{T}', Q_{\tau-1}, A'_{\tau-1})$ where $A'_{\tau-1}$ is the action taken by \mathcal{A}' at node $Q_{\tau-1}$. In this case, if $\text{children}(\mathcal{T}', Q_{\tau-1}, A'_{\tau-1}) = \emptyset$, then $Q_{\tau-1}$ would be fathomed by \mathcal{A}' , which cannot happen by the first case. Otherwise, if $\text{children}(\mathcal{T}', Q_{\tau-1}, A'_{\tau-1}) \neq \emptyset$, we show that we arrive at a contradiction due to the fact that the scoring rules, action-set functions, and children functions are all path-wise. Let $A'_{\tau-1}$ denote the action taken by \mathcal{A}' at $Q_{\tau-1}$, and let $A_{\tau-1}$ denote the action taken by Algorithm 1 at $Q_{\tau-1}$. Since actions is path-wise,

$$\text{actions}(\mathcal{T}, Q_{\tau-1}) = \text{actions}(\mathcal{T}_{Q_{\tau-1}}, Q_{\tau-1}) = \text{actions}(\mathcal{T}', Q_{\tau-1}).$$

Since ascore_1 and ascore_2 are path-wise, we have

$$\begin{aligned} \mu \cdot \text{ascore}_1(\mathcal{T}, Q_{\tau-1}, A) + (1 - \mu) \cdot \text{ascore}_2(\mathcal{T}, Q_{\tau-1}, A) \\ = \mu \cdot \text{ascore}_1(\mathcal{T}_{Q_{\tau-1}}, Q_{\tau-1}, A) + (1 - \mu) \cdot \text{ascore}_2(\mathcal{T}_{Q_{\tau-1}}, Q_{\tau-1}, A) \\ = \mu \cdot \text{ascore}_1(\mathcal{T}', Q_{\tau-1}, A) + (1 - \mu) \cdot \text{ascore}_2(\mathcal{T}', Q_{\tau-1}, A). \end{aligned}$$

for all actions $A \in \text{actions}(\mathcal{T}_{Q_{\tau-1}}, Q_{\tau-1})$. Therefore Algorithm 1 and \mathcal{A}' choose the same action at node $Q_{\tau-1}$, that is, $A_{\tau-1} = A'_{\tau-1}$. Finally, since children is path-wise, we have

$$\text{children}(\mathcal{T}, Q_{\tau-1}, A_{\tau-1}) = \text{children}(\mathcal{T}_{Q_{\tau-1}}, Q_{\tau-1}, A_{\tau-1}) = \text{children}(\mathcal{T}', Q_{\tau-1}, A_{\tau-1}).$$

Since $Q_\tau \in \text{children}(\mathcal{T}, Q_{\tau-1}, A_{\tau-1})$, this is a contradiction, which completes the proof. \blacktriangleleft

We use the following generalization of Claim 3.4 of Balcan et al. [7] that shows the behavior of \mathcal{A}' is piecewise constant. While their argument only applies to single-variable branching, our key insight is that the same reasoning can be readily adapted to handle any actions (including general branching constraints and cutting planes). The structure of our proof (which we defer to the appendix) is identical, but is modified to work in our more general setting. This style of analysis is similar in spirit to [34].

► **Lemma 3.** *Let ascore_1 and ascore_2 be two path-wise action-selection scores. Fix the input root node Q . There are $T \leq k^{\Delta(\Delta-1)/2} b^\Delta$ subintervals I_1, \dots, I_T partitioning $[0, 1]$ where for any subinterval I_t , the action-selection score $\mu \cdot \text{ascore}_1 + (1 - \mu) \cdot \text{ascore}_2$ results in the same tree built by \mathcal{A}' for all $\mu \in I_t$, no matter what node selection policy is used.*

We now prove our main structural result for Algorithm 1.

► **Lemma 4.** *Let ascore_1 and ascore_2 be path-wise action-selection scores and let nscore_1 and nscore_2 be path-wise node-selection scores. Fix the input root node Q . There are $T \leq k^{\Delta(9+\Delta)} b^\Delta$ rectangles partitioning $[0, 1]^2$ such that for any rectangle R_t , the node-selection score $\lambda \cdot \text{nscore}_1 + (1 - \lambda) \cdot \text{nscore}_2$ and the action-selection score $\mu \cdot \text{ascore}_1 + (1 - \mu) \cdot \text{ascore}_2$ result in the same tree built by Algorithm 1 for all $(\mu, \lambda) \in R_t$.*

Proof. By Lemma 3, there is a partition of $[0, 1]$ into subintervals $I_1 \cup \dots \cup I_T$ such that for all μ within a given subinterval, the tree built by \mathcal{A}' is invariant (independent of the node-selection score). Fix a subinterval I_t of this partition. Let \mathcal{T} denote the tree built by Algorithm 1. For each node $Q \in \mathcal{T}$, let \mathcal{T}_Q denote the path from the root to Q in \mathcal{T} . Since nscore_1 is path-wise, for any tree \mathcal{T}' containing \mathcal{T}_Q as a rooted path, $\text{nscore}_1(\mathcal{T}', Q) = \text{nscore}_1(\mathcal{T}_Q, Q)$. The same holds for nscore_2 . For every pair of nodes $Q_1, Q_2 \in \mathcal{T}$, let $\lambda(Q_1, Q_2) \in [0, 1]$ denote the unique solution to

$$\begin{aligned} \lambda \cdot \text{nscore}_1(\mathcal{T}_{Q_1}, Q_1) + (1 - \lambda) \cdot \text{nscore}_2(\mathcal{T}_{Q_1}, Q_1) \\ = \lambda \cdot \text{nscore}_1(\mathcal{T}_{Q_2}, Q_2) + (1 - \lambda) \cdot \text{nscore}_2(\mathcal{T}_{Q_2}, Q_2), \end{aligned}$$

if it exists (if there are either (1) no solutions or (2) infinitely many solutions, set $\lambda(Q_1, Q_2) = 0$). The thresholds $\lambda(Q_1, Q_2)$ for every pair of nodes $Q_1, Q_2 \in \mathcal{T}$ partition $[0, 1]$ into subintervals such that for all λ within a given subinterval, the total order over the nodes of \mathcal{T} induced by nscore_λ is invariant. In particular, this means that the node selected by each iteration of Algorithm 1 is invariant. Let $J_1 \cup \dots \cup J_S$ denote these subintervals induced by the thresholds over all subinterval $I_t \in \{I_1, \dots, I_T\}$ established in Lemma 3.

We now show that this implies that the tree built by Algorithm 1 is invariant over all (μ, λ) within a given rectangle $I_t \times J_s$. Fix some rectangle $I_t \times J_s$. We proceed by induction on the iterations (of the while loop) of Algorithm 1. For the base case (iteration 0, before entering the while loop), the tree consists of only the root, so the hypothesis trivially holds. Now, suppose the statement holds up until the j th iteration, for some j . We analyze each line of Algorithm 1 to show that the behavior of the $j+1$ st iteration is independent of $(\mu, \lambda) \in I_t \times J_s$. First, since J_s determines the node selected at each iteration (as argued above), the node selected on the $j+1$ st iteration (line 3) is fixed, independent of $(\mu, \lambda) \in I_t \times J_s$. Denote this node by Q . Thus, whether $\text{depth}(Q) = \Delta$ is independent of $(\mu, \lambda) \in I_t \times J_s$, and similarly whether $\text{fathom}(\mathcal{T}, Q, \text{None}) = \text{true}$ is independent of $(\mu, \lambda) \in I_t \times J_s$ (line 4). This implies that whether or not Q is fathomed at this stage is independent of $(\mu, \lambda) \in I_t \times J_s$. If Q was fathomed, we are done. Otherwise, we argue that the action selected at line 7 is invariant over $(\mu, \lambda) \in I_t \times J_s$. By Lemma 3, \mathcal{A}' builds the same tree for all $\mu \in I_t$. Let \mathcal{T}_Q denote the path from the root to Q in this tree. By Lemma 2, \mathcal{T}_Q is the path from the root to Q in the tree built by Algorithm 1 as well. The action selected at Q by \mathcal{A}' is invariant over $\mu \in I_t$ (by Lemma 3). Therefore, since actions , ascore_1 , and ascore_2 are path-wise, the action A selected by Algorithm 1 at Q is invariant over $\mu \in I_t$. Finally, $\text{fathom}(\mathcal{T}, Q, A)$ and $\text{children}(\mathcal{T}, Q, A)$ are completely determined, so the execution of the remaining conditional statement (line 8 to line 13) is invariant over $(\mu, \lambda) \in I_t \times J_s$. Thus, the entire iteration of Algorithm 1 is invariant over $(\mu, \lambda) \in I_t \times J_s$, which completes the induction.

Finally, we count the total number of rectangles in our partition of $[0, 1]^2$. For each interval I_t in the partition established in Lemma 3, we obtained a partition of $I_t \times [0, 1]$ into rectangles induced by at most $\binom{|T|}{2}$ thresholds, which consists of at most at most

$$1 + \binom{(k^{\Delta+1} - 1)/(k - 1)}{2} \leq 1 + \left(\frac{k^{\Delta+1} - 1}{k - 1} \right)^2 \leq k^{5\Delta}$$

subintervals. Accounting for every interval $I_t \in \{I_1, \dots, I_T\}$ in the partition from Lemma 3, we get a total of $Tk^{5\Delta} \leq k^{\Delta(9+\Delta)/2} b^{\Delta}$ rectangles, as desired. \blacktriangleleft

We now derive generalization guarantees for the collection $\mathcal{F} = \{\text{cost}_{\mu, \lambda} : (\mu, \lambda) \in [0, 1]^2\}$ where cost is any tree-constant function, such as tree size. We do this by bounding the *pseudo-dimension* of \mathcal{F} , which is a combinatorial measure of intrinsic complexity of a class of

real valued functions. The pseudo-dimension of \mathcal{F} , denoted by $\text{Pdim}(\mathcal{F})$, is the largest positive integer N such that there exist N nodes $Q_1, \dots, Q_N \in \mathcal{Q}$ and N thresholds $r_1, \dots, r_N \in \mathbb{R}$ such that $|\{(\text{sign}(f(Q_1) - r_1), \dots, \text{sign}(f(Q_N) - r_N)) : f \in \mathcal{F}\}| = 2^N$. A well-known result in learning theory [3] states that if functions in \mathcal{F} have bounded range $[-H, H]$, then

$$N_{\mathcal{F}}(\varepsilon, \delta) = O\left(\frac{H^2}{\varepsilon^2} (\text{Pdim}(\mathcal{F}) + \ln(1/\delta))\right) \text{ and } \varepsilon_{\mathcal{F}}(N, \delta) = O\left(H \sqrt{\frac{\text{Pdim}(\mathcal{F}) + \ln(1/\delta)}{N}}\right).$$

When each function in \mathcal{F} maps to $\{0, 1\}$, the pseudo-dimension is more commonly referred to as the *VC dimension*.

Bounding the pseudo-dimension is a simple instantiation of the general framework provided by Balcan et al. [6] with the piecewise structure established in Lemma 4. Balcan et al.'s [6] main result gives pseudo-dimension bounds for families of piecewise structured functions in terms of the VC dimension of the class of 0/1 classifiers defining the boundaries of the functions, the number of classifiers defining the boundaries, and the pseudo-dimension of the family of functions when restricted to each piece. (Strictly, this result is in terms of the dual classes of the boundary and piece functions. However, since the dual class of all linear separators is the set of all linear separators, we omit this detail for simplicity.)

► **Theorem 5.** *Let $\text{cost}(Q)$ be any tree-constant cost function, and let $\text{cost}_{\mu, \lambda}(Q)$ be the cost of the tree built by Algorithm 1 on input root node Q using action-selection score parameterized by μ and node-selection score parameterized by λ . Then, $\text{Pdim}(\{\text{cost}_{\mu, \lambda}\}) = O(\Delta^2 \log k + \Delta \log b)$.*

Proof. By Lemma 4, there are at most $T = k^{\Delta(9+\Delta)}b^{\Delta}$ rectangles partitioning $[0, 1]^2$ such that for a fixed input node Q , $\text{cost}_{\mu, \lambda}(Q)$ is constant over each rectangle as a function of μ, λ . These T rectangles can be defined by T thresholds on $[0, 1]$ corresponding to μ and T thresholds on $[0, 1]$ corresponding to λ . Thus, the T rectangles can be identified by $T^2 = k^{2\Delta(9+\Delta)}b^{2\Delta}$ linear separators in \mathbb{R}^2 . The VC dimension of linear separators in \mathbb{R}^2 is $O(1)$. The pseudo-dimension of the set of constant functions is also $O(1)$. Plugging these quantities into the main theorem of Balcan et al. [6] yields the theorem statement. ◀

3.1 Multiple actions

Theorem 5 can be easily generalized to the case where there are multiple actions of different types taken at each node of Algorithm 1. Specifically, there are now d path-wise action-set functions $\text{actions}_1, \dots, \text{actions}_d$, and at line 7 of Algorithm 1 we take one action of each type, that is, we select action $A_1 \in \text{actions}_1(\mathcal{T}, Q)$, $A_2 \in \text{actions}_2(\mathcal{T}, Q)$, and so on. The functions fathom and children then depend on all d actions taken at node Q . We assume that there are two scoring rules ascore_1^i and ascore_2^i for each action type $i = 1, \dots, d$. Algorithm 1 can then be parameterized by (μ, λ) , where $\mu \in \mathbb{R}^d$ is a vector of parameters controlling each action, so the i th action is selected to maximize $\mu_i \cdot \text{ascore}_1^i + (1 - \mu_i) \cdot \text{ascore}_2^i$. Then, as long as $d = O(1)$, we get the same pseudo-dimension bound. We assume b is a uniform upper bound on the size of actions_i for any i . The proof is nearly identical, and we defer it to the appendix (which also contains more details on the multiple-action setup).

► **Theorem 6.** *Let $\text{cost}(Q)$ be any tree-constant cost function, and let $\text{cost}_{\mu, \lambda}(Q)$ be the cost of the tree built by Algorithm 1 on input root node Q using action-selection scores parameterized by $\mu \in \mathbb{R}^d$, where $d = O(1)$, and node-selection score parameterized by λ . Then, $\text{Pdim}(\{\text{cost}_{\mu, \lambda}\}) = O(\Delta^2 \log k + \Delta \log b)$.*

4 Branch-and-cut for integer programming

We now instantiate our main results with the three main components of the B&C algorithm: branching, cutting planes, and node selection, used to solve IPs $\max\{c^T x : Ax \leq b, x \geq 0, x \in \mathbb{Z}^n\}$ where $c \in \mathbb{R}^n$, $A \in \mathbb{Z}^{m \times n}$, $b \in \mathbb{Z}^m$. The function `fathom`(\mathcal{T}, Q, A) outputs `true` if after having taken action A the LP relaxation at Q is integral, infeasible, or worse than the best integral solution found so far in \mathcal{T} . The function `children`(\mathcal{T}, Q, A) outputs the two subproblems generated by the branching procedure on the IP at Q after having taken action A . For simplicity we refer only to IPs, but everything in our discussion applies to mixed IPs as well. In our model of tree search, node selection is controlled by λ . Cutting planes and branching are types of actions and controlled by μ .

4.1 Branching

In this section, we provide guarantees for branching. Throughout this section we assume $\Delta = O(n)$, as is the case with single-variable branching.

4.1.1 Multivariable branching constraints

It is well known that allowing for more general generation of branching constraints can result in smaller B&C trees. Gilpin and Sandholm [17] studied multivariable branches of the form $\sum_{i \in S} x[i] \leq \lfloor \sum_{i \in S} x_{LP}^*[i] \rfloor$, $\sum_{i \in S} x[i] \geq \lceil \sum_{i \in S} x_{LP}^*[i] \rceil$ where S is a subset of the integer variables such that $\sum_{i \in S} x_{LP}^*[i] \notin \mathbb{Z}$. Here, $\text{actions}(\mathcal{T}, Q) = 2^{\lfloor n \rfloor}$, so, $\text{Pdim}(\{\text{cost}_{\mu, \lambda}\}) = O(n^2)$. So our sample complexity bound for multivariable branching constraints is, surprisingly, only a constant factor worse than the bound for single-variable branching constraints.

We give a simple example where B&C using only single variable branches builds a tree of exponential size, while a single branch on the entire set of variables at the root yields two infeasible subproblems (and a B&C tree of size 3).

► **Theorem 7.** *For any n , there is an IP with two constraints and n variables such that with only single variable branches, B&C builds a tree of size $2^{(n-1)/2}$, while with a suitable multivariable branch, B&C builds a tree of size three.*

Proof. Let n be an odd positive integer. Consider the infeasible IP $\max\{\sum_{i=1}^n x[i] : 2 \sum_{i=1}^n x[i] = n, x \in \{0, 1\}^n\}$. Jeroslow [26] proved that with only single-variable branches, B&C builds a tree with $2^{(n-1)/2}$ nodes to determine infeasibility. However, with a suitable multivariable branch, B&C will build a tree of constant size. The optimal solution to the LP relaxation of the IP is attained when all variables are set to $1/2$. A multivariable branch on all n variables produces the two subproblems with constraints $\sum_{i=1}^n x[i] \leq \lfloor n/2 \rfloor$ and $\sum_{i=1}^n x[i] \geq \lceil n/2 \rceil$, respectively. Since n is odd, $\lfloor n/2 \rfloor < n/2$ and $\lceil n/2 \rceil > n/2$, so the LP relaxations of both subproblems are infeasible. Thus, B&C builds a tree with three nodes. ◀

Yang et al. [45] provide more examples of situations where multivariable branching yields dramatic improvements in tree size over single variable branching. They also perform a computational evaluation of a few different strategies for generating multivariable branching constraints. Yang et al. [44] explore gradient-boosting for learning to mimic strong branching for multiple variables.

4.1.2 Branching on general disjunctions

Branching constraints can be even more general than multivariable branches. Given any integer vector $\pi \in \mathbb{Z}^n$ and any integer $\pi_0 \in \mathbb{Z}$ (jointly referred to as a *disjunction*), the constraints $\pi^T x \leq \pi_0$ or $\pi^T x \geq \pi_0 + 1$ represent a valid partition of the feasible region

into subproblems. Owen and Mehrotra [35] ran the first experiments demonstrating that branching on general disjunctions can lead to significantly smaller tree sizes. Subsequent works have posed different heuristics to select disjunctions to branch on [14, 33].

In practice it is known that additional IP constraints should not have coefficients that are too large. If C is a bound on the magnitude of the coefficient of any disjunction, then $\text{actions}(\mathcal{T}, Q) = \{-C, \dots, C\}^{n+1}$, so $\text{Pdim}(\{\text{cost}_{\mu, \lambda}\}) = O(n^2 \log C)$. Karamanov and Cornuéjols [28] conduct a computational evaluation of disjunctions derived from Gomory mixed-integer cuts. In this setting, $\text{actions}(\mathcal{T}, Q)$ is the set of m or fewer disjunctions corresponding to the m or fewer Gomory mixed-integer cuts derived from the simplex tableau from solving the LP relaxation of Q . In this case, $\text{Pdim}(\{\text{cost}_{\mu, \lambda}\}) = O(n^2 + n \log m)$.

4.2 Cutting planes


The action set can also correspond to cutting planes used to refine the feasible region of the IP at any stage of B&C. Here, $\text{actions}(\mathcal{T}, Q)$ is any set of cutting planes derived solely using the path from the root to the IP at Q . Examples include the set of Chvátal-Gomory (CG) derived from the simplex tableau [18], and various combinatorial families of cutting planes such as clique cuts, odd-hole cuts, and cover cuts. The set $\text{actions}(\mathcal{T}, Q)$ can also consist of sequences of cutting planes, representing adding several cutting planes to the IP in waves. For example, the set of all sequences of w CG cuts generated from the simplex tableau for an IP with m constraints has size at most m^w (regardless of whether the LP is resolved after each cut). The number of such cutting planes provided by the LP tableau at any node in the tree is at most $O(m + nw)$ (the original IP has m constraints, and after at most n branches there are an additional n branching constraints and at most nw cutting planes), which means that $|\text{actions}(\mathcal{T}, Q)| \leq O(m + nw)^w$. Thus, $\text{Pdim}(\{\text{cost}_{\mu, \lambda}\}) = O(n^2 + nw \log(m + nw))$.

We can also handle arbitrary CG cuts (not just ones from the LP tableau). Balcan et al. [9] proved that given an IP with feasible region $\{\mathbf{x} \in \mathbb{Z}^n : A\mathbf{x} \leq \mathbf{b}, \mathbf{x} \geq 0\}$, even though there are infinitely many CG cut parameters, there are effectively only $O(w2^w \|A\|_{1,1} + 2^w \|\mathbf{b}\|_1 + nw)^{1+mw}$ distinct sequences of cutting planes that w CG cut parameters can produce. At any node in the B&C tree, the number of constraints is at most $O(m + nw)$. So, on the domain of IPs with $\|A\|_{1,1} \leq \alpha$ and $\|\mathbf{b}\|_1 \leq \beta$, $|\text{actions}(\mathcal{T}, Q)| \leq O(w2^w \alpha + 2^w \beta + nw)^{1+w \cdot O(m+nw)}$. Thus, $\text{Pdim}(\{\text{cost}_{\mu, \lambda}\}) = O(n^2 w^3 m \log(\alpha + \beta + n))$.


4.2.1 Experiments on cover cuts for the multiple knapsack problem

In this section, we demonstrate via experiments that tuning a convex combination of scoring rules to select cuts can lead to dramatically smaller branch-and-cut trees when done in a data-dependent manner. We study the classical NP-hard *multiple knapsack problem*: given a set N of items where each item $i \in N$ has a value $p_i \geq 0$ and a weight $w_i \geq 0$, and a set K of knapsacks where each knapsack $k \in K$ has a capacity $W_k \geq 0$, the goal is to find a feasible packing of the items into the knapsacks of maximum value. We assume, without loss of generality, that the items are labeled in descending order of weight, that is, $w_1 \geq w_2 \geq \dots \geq w_{|N|}$. This problem can be formulated as the following binary IP:

$$\begin{aligned} & \text{maximize} && \sum_{i \in N} \sum_{k \in K} p_i x_{k,i} \\ & \text{subject to} && \sum_{i \in N} w_i x_{k,i} \leq W_k && \forall k \in K \\ & && \sum_{k \in K} x_{k,i} \leq 1 && \forall i \in N \\ & && x_{k,i} \in \{0, 1\} && \forall i \in N, k \in K \end{aligned}$$



■ **Figure 1** Chvátal distribution with 35 items and 2 knapsacks.




■ **Figure 2** Chvátal distribution with 35 items and 3 knapsacks.


A subset $C \subseteq N$ of items is called a *cover* for knapsack $k \in K$ if $\sum_{i \in C} w_i > W_k$. If C is a cover, no feasible solution can have $x_{k,i} = 1$ for all $i \in C$, so $\sum_{i \in C} x_{k,i} \leq |C| - 1$ is a valid constraint – called a *cover cut*. When C is minimal (that is, $C \setminus \{i\}$ is not a cover for every $i \in C$), such cover cuts help tighten the knapsack IP by cutting off fractional LP solutions. We generate (a subset of all) cover cuts for each knapsack k as follows: for each $i \in N$, let $j > i$ be minimal such that $C = \{i, i + 1, \dots, j\}$ is a cover for k (if such a j exists). Since $w_i \geq w_j$ for $j > i$, C is a minimal cover, and moreover the *extended cover cut* $\sum_{i=1}^j x_i \leq |C| - 1$ is valid and dominates the minimal cover cut $\sum_{i \in C} x_i \leq |C| - 1$. Extended cover cuts generated from minimal covers are known to be facet defining for the integer hull under certain natural conditions [12].

We investigate the relationship between three scoring rules for cutting planes. The first is *efficacy* (E), which is the perpendicular distance from the current LP solution to the cutting plane. The second is *parallelism* (P), which measures the angle between the objective and the normal vector to the cutting plane. The third is *directed cutoff* (D), which is the distance from the current LP solution to the cutting plane along the direction of the line segment connecting the LP solution to the current best incumbent integer solution. More details, including explicit formulas, can be found in [9] and references therein.

We consider two specific instances of the multiple knapsack problem, which are loosely based on a class of knapsack problems introduced by Chvátal that are difficult to solve with vanilla branch-and-bound [11, 45]. In the first, $p_i = w_i$ for all $i \in N$, and $W_k = \lfloor (\sum_{i \in N} w_i) / 2|K| \rfloor + (k - 1)$ for each $k = 1, \dots, |K|$. In the second, $p_i = w_{|N|-i+1}$, so the most valuable item is the lightest and the least valuable item is the heaviest, and W_k is defined as in the first type. We call the first class of problems *Chvátal instances* and the second class *reverse Chvátal instances*. For a given N, K , we generate (reverse) Chvátal instances by drawing each weight independently as $w_i = \lfloor z_i \rfloor$, where $z_i \sim \mathcal{N}(50, 2)$, and sorting the items by weight in descending order.



■ **Figure 3** Reverse Chvátal distribution with 100 items and 10 knapsacks.



■ **Figure 4** Reverse Chvátal distribution with 100 items and 15 knapsacks.

In our experiments, we add (whenever possible) two extended cover cuts obtained in the aforementioned manner at every node of the B&C tree. The two cuts chosen are the two with the highest score $\mu \cdot \text{ascore}_1 + (1 - \mu) \cdot \text{ascore}_2$ among all extended cover cuts that are violated by the current LP optimum, where $\text{ascore}_1, \text{ascore}_2 \in \{E, D, P\}$. Figures 1-4 display the average tree size over 1000 samples for different Chvátal and reverse Chvátal distributions as a function of μ , where the domain $[0, 1]$ of μ is discretized in increments of 0.01. We ran our experiments using the Python API of CPLEX 12.10 with default cut generation turned off. All other aspects of B&C (e.g. variable and node selection) are controlled by the default settings of CPLEX. The key takeaway of our plots is that tuning a convex combination of scoring rules can lead to significant savings in B&C tree size, and that this tuning must be done with the IP distribution in mind. No single parameter produces small trees for all the distributions we considered, and in fact a μ that minimizes tree size for one distribution can result in the largest trees for another (as in Figures 2b and 4b, for example). Furthermore, many of the plots display discernible trends (and in some cases are quite smooth), suggesting that the number of samples required to avoid overfitting in practice can be significantly smaller than our theoretical bounds.

4.3 Improved bounds for branch-and-cut

To allow node selection, branching, and cutting-plane selection to be tuned simultaneously, we apply Theorem 6. This allows us to bound the pseudo-dimension of the family of functions $\{\text{cost}_{\mu_1, \mu_2, \lambda}\}$, where μ_1 controls branching, μ_2 controls cutting-plane selection, and λ controls node selection. Let $\text{actions}_1(\mathcal{T}, Q)$ denote the set of branching actions available at Q , and let $\text{actions}_2(\mathcal{T}, Q)$ denote the set of cutting planes available at Q . Let $b_1, b_2 \in \mathbb{N}$ be such that $|\text{actions}_1(\mathcal{T}, Q)| \leq b_1$ and $|\text{actions}_2(\mathcal{T}, Q)| \leq b_2$ for all \mathcal{T} and all $Q \in \mathcal{T}$. Fix two branching scores $\text{ascore}_1^1, \text{ascore}_1^2$, fix two cutting-plane selection scores $\text{ascore}_2^1, \text{ascore}_2^2$, and fix two node-selection scores $\text{nscore}_1, \text{nscore}_2$.

► **Theorem 8.** Let $\text{cost}(Q)$ be any tree-constant cost function, and let $\text{cost}_{\mu_1, \mu_2, \lambda}$ be the cost of the tree built by B&C using branching score $\mu_1 \cdot \text{ascore}_1^1 + (1 - \mu_1) \cdot \text{ascore}_2^1$, cutting-plane selection score $\mu_2 \cdot \text{ascore}_1^2 + (1 - \mu_2) \cdot \text{ascore}_2^2$, and node-selection score $\lambda \cdot \text{nscore}_1 + (1 - \lambda) \cdot \text{nscore}_2$. Then, with $\Delta = O(n)$, $\text{Pdim}(\{\text{cost}_{\mu_1, \mu_2, \lambda}\}) = O(n^2 + n \log(b_1 + b_2))$.

4.3.1 Comparison to existing bounds

Balcan et al. [9] give a pseudo-dimension bound for tree search with a linear dependence on a cap κ on the number of nodes allowed in any tree. Their pseudo-dimension bound in our setting is $\text{Pdim}(\{\text{cost}_{\mu_1, \mu_2, \lambda}\}) = O(\kappa \log \kappa + \kappa \log b_1 + \kappa \log b_2)$. While κ is treated as a constant, it can be a prohibitively large quantity. In fact, without explicitly enforcing a limit on the number of nodes expanded by B&C, Balcan et al. [9] obtain a pseudo-dimension bound of $O(2^n (\log b_1 + \log b_2))$. Balcan et al. [7] use the path-wise property to prove that $\text{Pdim}(\{\text{cost}_\mu\}) = O(n^2)$ for single-variable branching, but for the case where branching is the only tunable component of B&C (and node selection is fixed).

5 Conclusions and future research

We presented a general model of tree search and proved sample complexity guarantees for this model that improve and generalize upon the recent sample complexity theory for configuring branch-and-cut. There are many interesting and open directions for future research. One compelling open question is to obtain pseudo-dimension bounds when action sets are infinite. Balcan et al. [9] alluded to this question in the case of cutting planes, and neither the techniques of their work nor the techniques of the present work can handle, for example, important infinite cutting-plane families such as the class of Gomory mixed-integer cuts, or the infinitely many valid disjunctions that could be branched on. Beyond integer programming, our model of tree search could potentially be applied to completely different problem domains that exhibit tree structure. Another direction is to extend our results to convex combinations of $\ell > 2$ scoring rules $\mu_1 \text{score}_1 + \dots \mu_\ell \text{score}_\ell$, as Balcan et al. [7] do in the special case of single-variable branching. However, their pseudo-dimension bound grows exponentially in the number of variables n in that special case; developing techniques that lead to a polynomial dependence on n remains a challenging open question.

References

- 1 Tobias Achterberg. *Constraint Integer Programming*. PhD thesis, Technische Universität Berlin, 2007.
- 2 Alejandro Marcos Alvarez, Quentin Louveaux, and Louis Wehenkel. A machine learning-based approximation of strong branching. *INFORMS Journal on Computing*, 29(1):185–195, 2017.
- 3 Martin Anthony and Peter Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, 2009.
- 4 Egon Balas, Sebastián Ceria, and Gérard Cornuéjols. Mixed 0-1 programming by lift-and-project in a branch-and-cut framework. *Management Science*, 42(9):1229–1246, 1996.
- 5 Maria-Florina Balcan. Data-driven algorithm design. In Tim Roughgarden, editor, *Beyond Worst Case Analysis of Algorithms*. Cambridge University Press, 2020.
- 6 Maria-Florina Balcan, Dan DeBlasio, Travis Dick, Carl Kingsford, Tuomas Sandholm, and Ellen Vitercik. How much data is sufficient to learn high-performing algorithms? Generalization guarantees for data-driven algorithm design. In *Annual Symposium on Theory of Computing (STOC)*, 2021.

- 7 Maria-Florina Balcan, Travis Dick, Tuomas Sandholm, and Ellen Vitercik. Learning to branch. In *International Conference on Machine Learning (ICML)*, 2018.
- 8 Maria-Florina Balcan, Vaishnavh Nagarajan, Ellen Vitercik, and Colin White. Learning-theoretic foundations of algorithm configuration for combinatorial partitioning problems. In *Conference on Learning Theory (COLT)*, 2017.
- 9 Maria-Florina Balcan, Siddharth Prasad, Tuomas Sandholm, and Ellen Vitercik. Sample complexity of tree search configuration: Cutting planes and beyond. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- 10 Antonia Chmiela, Elias B Khalil, Ambros Gleixner, Andrea Lodi, and Sebastian Pokutta. Learning to schedule heuristics in branch-and-bound. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- 11 Vasek Chvátal. Hard knapsack problems. *Operations Research*, 28(6):1402–1411, 1980.
- 12 Michele Conforti, Gérard Cornuéjols, Giacomo Zambelli, et al. *Integer programming*, volume 271. Springer, 2014.
- 13 Giovanni Di Liberto, Serdar Kadioglu, Kevin Leo, and Yuri Malitsky. Dash: Dynamic approach for switching heuristics. *European Journal of Operational Research*, 248(3):943–953, 2016.
- 14 Matteo Fischetti and Andrea Lodi. Local branching. *Mathematical Programming*, 98:23–47, 2002.
- 15 Gerald Gamrath, Daniel Anderson, Ksenia Bestuzheva, Wei-Kun Chen, Leon Eifler, Maxime Gasse, Patrick Gemander, Ambros Gleixner, Leona Gottwald, Katrin Halbig, Gregor Hendel, Christopher Hojny, Thorsten Koch, Pierre Le Bodic, Stephen J. Maher, Frederic Matter, Matthias Miltenberger, Erik Mühmer, Benjamin Müller, Marc E. Pfetsch, Franziska Schlösser, Felipe Serrano, Yuji Shinano, Christine Tawfik, Stefan Vigerske, Fabian Wegscheider, Dieter Weninger, and Jakob Witzig. The SCIP Optimization Suite 7.0. Technical report, Optimization Online, March 2020. URL: http://www.optimization-online.org/DB_HTML/2020/03/7705.html.
- 16 Maxime Gasse, Didier Chételat, Nicola Ferroni, Laurent Charlin, and Andrea Lodi. Exact combinatorial optimization with graph convolutional neural networks. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 15554–15566, 2019.
- 17 Andrew Gilpin and Tuomas Sandholm. Information-theoretic approaches to branching in search. *Discrete Optimization*, 8(2):147–159, 2011. Early version in IJCAI-07.
- 18 Ralph E. Gomory. Outline of an algorithm for integer solutions to linear programs. *Bulletin of the American Mathematical Society*, 64(5):275–278, 1958.
- 19 Prateek Gupta, Maxime Gasse, Elias B Khalil, M Pawan Kumar, Andrea Lodi, and Yoshua Bengio. Hybrid models for learning to branch. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- 20 Rishi Gupta and Tim Roughgarden. A PAC approach to application-specific algorithm selection. *SIAM Journal on Computing*, 46(3):992–1017, 2017.
- 21 He He, Hal Daume III, and Jason M Eisner. Learning to search in branch and bound algorithms. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2014.
- 22 Eric Horvitz, Yongshao Ruan, Carla Gomez, Henry Kautz, Bart Selman, and Max Chickering. A Bayesian approach to tackling hard computational problems. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2001.
- 23 Zeren Huang, Kerong Wang, Furui Liu, Hui-ling Zhen, Weinan Zhang, Mingxuan Yuan, Jianye Hao, Yong Yu, and Jun Wang. Learning to select cuts for efficient mixed-integer programming. *arXiv preprint*, 2021. [arXiv:2105.13645](https://arxiv.org/abs/2105.13645).
- 24 Frank Hutter, Holger Hoos, and Kevin Leyton-Brown. Sequential model-based optimization for general algorithm configuration. In *International Conference on Learning and Intelligent Optimization (LION)*, pages 507–523, 2011.
- 25 Frank Hutter, Holger H Hoos, Kevin Leyton-Brown, and Thomas Stützle. ParamILS: An automatic algorithm configuration framework. *Journal of Artificial Intelligence Research*, 36(1):267–306, 2009.

- 26 Robert G Jeroslow. Trivial integer programs unsolvable by branch-and-bound. *Mathematical Programming*, 6(1):105–109, 1974.
- 27 Serdar Kadioglu, Yuri Malitsky, Meinolf Sellmann, and Kevin Tierney. ISAC—instance-specific algorithm configuration. In *European Conference on Artificial Intelligence (ECAI)*, 2010.
- 28 Miroslav Karamanov and Gérard Cornuéjols. Branching on general disjunctions. *Mathematical Programming*, 128(1-2):403–436, 2011.
- 29 Elias Khalil, Pierre Le Bodic, Le Song, George Nemhauser, and Bistra Dilkina. Learning to branch in mixed integer programming. In *AAAI Conference on Artificial Intelligence*, 2016.
- 30 Elias Khalil, Bistra Dilkina, George Nemhauser, Shabbir Ahmed, and Yufen Shao. Learning to run heuristics in tree search. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2017.
- 31 Robert Kleinberg, Kevin Leyton-Brown, and Brendan Lucier. Efficiency through procrastination: Approximately optimal algorithm configuration with runtime guarantees. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2017.
- 32 Kevin Leyton-Brown, Eugene Nudelman, and Yoav Shoham. Empirical hardness models: Methodology and a case study on combinatorial auctions. *Journal of the ACM*, 56(4):1–52, 2009.
- 33 Ashutosh Mahajan and Theodore K Ralphs. Experiments with branching using general disjunctions. In *Operations Research and Cyber-Infrastructure*, pages 101–118. Springer, 2009.
- 34 Nimrod Megiddo. Combinatorial optimization with rational objective functions. *Mathematics of Operations Research*, pages 414–424, 1979.
- 35 Jonathan H. Owen and Sanjay Mehrotra. Experimental results on using general disjunctions in branch-and-bound for general-integer linear programs. *Computational Optimization and Applications*, 20(2):159–170, November 2001.
- 36 Ashish Sabharwal, Horst Samulowitz, and Chandra Reddy. Guiding combinatorial optimization with UCT. In *International Conference on AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*. Springer, 2012.
- 37 Tuomas Sandholm. Very-large-scale generalized combinatorial multi-attribute auctions: Lessons from conducting \$60 billion of sourcing. In Zvika Neeman, Alvin Roth, and Nir Vulkan, editors, *Handbook of Market Design*. Oxford University Press, 2013.
- 38 Jialin Song, Ravi Lanka, Yisong Yue, and Bistra Dilkina. A general large neighborhood search framework for solving integer programs. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- 39 Yunhao Tang, Shipra Agrawal, and Yuri Faenza. Reinforcement learning for integer programming: Learning to cut. *International Conference on Machine Learning (ICML)*, 2020.
- 40 Gellért Weisz, András György, and Csaba Szepesvári. LEAPSANDBOUNDS: A method for approximately optimal algorithm configuration. In *International Conference on Machine Learning (ICML)*, 2018.
- 41 Franz Wesselmann and Uwe Suhl. Implementing cutting plane management and selection techniques. Technical report, University of Paderborn, 2012.
- 42 Lin Xu, Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. Satzilla: portfolio-based algorithm selection for SAT. *Journal of Artificial Intelligence Research*, 32(1):565–606, 2008.
- 43 Lin Xu, Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. Hydra-MIP: Automated algorithm configuration and selection for mixed integer programming. In *RCRA workshop on Experimental Evaluation of Algorithms for Solving Problems with Combinatorial Explosion at the International Joint Conference on Artificial Intelligence (IJCAI)*, 2011.
- 44 Yu Yang, Natasha Boland, Bistra Dilkina, and Martin Savelsbergh. Learning generalized strong branching for set covering, set packing, and 0-1 knapsack problems. Technical report, Technical Report, 2020., 2020.
- 45 Yu Yang, Natasha Boland, and Martin Savelsbergh. Multivariable branching: A 0-1 knapsack problem case study. *INFORMS Journal on Computing*, 2021.

A Analysis of \mathcal{A}'

Proof of Lemma 3. Let \mathcal{T} denote the tree built by \mathcal{A}' . For $i \in [\Delta]$, let $\mathcal{T}[i]$ denote the restriction of \mathcal{T} to nodes of depth at most i . Let $\text{ascore}_\mu = \mu \cdot \text{ascore}_1 + (1 - \mu) \cdot \text{ascore}_2$. We prove the lemma by induction on i . In particular, we show that for each $i \in [\Delta]$, there are $k^{i(i-1)/2}b^i$ subintervals partitioning $[0, 1]$ such that $\mathcal{T}[i]$ is invariant over all μ within any given subinterval. Since $\mathcal{T}[\Delta] = \mathcal{T}$, this implies the lemma statement. The base case of $i = 1$ is trivial since $\mathcal{T}[1]$ consists of only the root.

Now, suppose the statement holds for some $i \in \{1, \dots, \Delta - 1\}$. That is, there are $T \leq k^{i(i-1)/2}b^i$ disjoint intervals $I_1 \cup \dots \cup I_T = [0, 1]$ such that $\mathcal{T}[i]$ is invariant over all μ within any given subinterval (our inductive hypothesis). Fix one of these subintervals I_t . We subdivide I_t into subintervals such that $\mathcal{T}[i+1]$ is invariant within each one of these smaller subintervals. Let Q be any leaf of $\mathcal{T}[i]$, and for $\mu \in I_t$ let \mathcal{T}_μ denote the state of the tree using ascore_μ at the point that Q is selected. Since $i < \Delta$, Q is not fathomed at line 4, regardless of μ . Next, since **actions** is path-wise, the actions available at Q depend only on the path \mathcal{T}_Q from the root of \mathcal{T} to Q , which, by the inductive hypothesis, is invariant over all $\mu \in I_t$. That is, $\text{actions}(\mathcal{T}_\mu, Q) = \text{actions}(\mathcal{T}_Q, Q)$ for all $\mu \in I_t$. Then, ascore_μ with parameter μ will select action $A \in \text{actions}(\mathcal{T}_Q, Q)$ if and only if

$$\begin{aligned} A &= \operatorname{argmax}_{A_0 \in \text{actions}(\mathcal{T}_Q, Q)} \mu \cdot \text{ascore}_1(\mathcal{T}_\mu, Q, A_0) + (1 - \mu) \cdot \text{ascore}_2(\mathcal{T}_\mu, Q, A_0) \\ &= \operatorname{argmax}_{A_0 \in \text{actions}(\mathcal{T}_Q, Q)} \mu \cdot \text{ascore}_1(\mathcal{T}_Q, Q, A_0) + (1 - \mu) \cdot \text{ascore}_2(\mathcal{T}_Q, Q, A_0), \end{aligned}$$

where the second equality follows from the fact that ascore_1 and ascore_2 are path-wise. Thus, for a fixed A_0 , ascore_μ is linear in μ , so for each A_0 there is at most one subinterval of $[0, 1]$ such that for all μ in that subinterval, A_0 maximizes ascore_μ . Thus, each leaf of $\mathcal{T}[i]$ contributes at most b subintervals such that for μ within a given subinterval, the action selected at each leaf of $\mathcal{T}[i]$ is invariant. $\mathcal{T}[i]$ consists of at most k^i leaves, so this is a total of at most $k^i b$ subintervals. Now, since the action A selected at each leaf Q of $\mathcal{T}[i]$ is invariant, the set of children $\text{children}(\mathcal{T}_\mu, Q, A) = \text{children}(\mathcal{T}_Q, Q, A)$ of Q added to the tree is also invariant, using the fact that **children** is path-wise. This shows that within every subinterval, $\mathcal{T}[i+1]$ is invariant. The total number of subintervals is, by the induction hypothesis, at most $k^{i(i-1)/2}b^i \cdot k^i b = k^{(i+1)i/2}b^{i+1}$, as desired. \blacktriangleleft

B Multiple actions

Let $\text{actions}_1, \dots, \text{actions}_d$ be path-wise. The multi-action version of Algorithm 1 is given by Algorithm 2. There are two scoring rules ascore_1^i and ascore_2^i for each action type $i \in [d]$. Algorithm 2 can then be parameterized by (μ, λ) , where $\mu \in \mathbb{R}^d$ is a vector of parameters controlling each action: the i th action is selected to maximize $\mu_i \cdot \text{ascore}_1^i + (1 - \mu_i) \cdot \text{ascore}_2^i$. As before, we assume there are $b, k \in \mathbb{N}$ such that $|\text{actions}_i(\mathcal{T}, Q)| \leq b$ for any i and any $Q \in \mathcal{T}$, and $|\text{children}(\mathcal{T}, Q, A_1, \dots, A_d)| \leq k$ for all Q, A_1, \dots, A_d .

Let \mathcal{A}' , as in the single-action setting, be Algorithm 2 with the evaluations of **fathom** on line 4 and line 8 suppressed. Then, we may prove a slight generalization of lemma 3.

► **Lemma 9.** *Let ascore_1^i and ascore_2^i be two path-wise action-selection scores, for each $i \in \{1, \dots, d\}$. Fix the input root node Q . There are $T \leq k^{d\Delta(\Delta-1)/2}b^{d\Delta}$ boxes of the form $R_t = I_1 \times \dots \times I_d$ partitioning $[0, 1]^d$ where for any box R_t , the action-selection scores $\mu_i \cdot \text{ascore}_1^i + (1 - \mu_i) \cdot \text{ascore}_2^i$ results in the same tree built by \mathcal{A}' for all $\mu \in R_t$, no matter what node selection policy is used.*

■ **Algorithm 2** Tree search with multiple actions.

Input: Root node Q , depth limit Δ

```

1: Initialize  $\mathcal{T} = Q$ .
2: while  $\mathcal{T}$  contains an unfathomed leaf do
3:   Select a leaf  $Q$  of  $\mathcal{T}$  that maximizes  $\text{nscore}(\mathcal{T}, Q)$ .
4:   if  $\text{depth}(Q) = \Delta$  or  $\text{fathom}(\mathcal{T}, Q, \text{None}, \dots, \text{None})$  then
5:     Fathom  $Q$ .
6:   else
7:     For  $i = 1, \dots, d$ , select  $A_i \in \text{actions}_i(\mathcal{T}, Q)$  that maximizes  $\text{ascore}_i(\mathcal{T}, Q, A_i)$ .
8:     if  $\text{fathom}(\mathcal{T}, Q, A_1, \dots, A_d)$  then
9:       Fathom  $Q$ .
10:    else if  $\text{children}(\mathcal{T}, Q, A_1, \dots, A_d) = \emptyset$  then
11:      Fathom  $Q$ .
12:    else
13:      Add all nodes in  $\text{children}(\mathcal{T}, Q, A_1, \dots, A_d)$  to  $\mathcal{T}$  as children of  $Q$ .
```

Proof. Let \mathcal{T} denote the tree built by \mathcal{A}' . For $i \in [\Delta]$, let $\mathcal{T}[i]$ denote the restriction of \mathcal{T} to nodes of depth at most i . Let $\text{ascore}_{\mu_i}^i = \mu_i \cdot \text{ascore}_1^i + (1 - \mu_i) \cdot \text{ascore}_2^i$. We prove the lemma by induction on i . In particular, we show that for each $i \in [\Delta]$, there are $k^{di(i-1)/2} b^{di}$ boxes partitioning $[0, 1]^d$ such that $\mathcal{T}[i]$ is invariant over all μ within any given box. Since $\mathcal{T}[\Delta] = \mathcal{T}$, this implies the lemma statement. The base case of $i = 1$ is trivial since $\mathcal{T}[1]$ consists of only the root, regardless of $\mu \in [0, 1]^d$.

Now, suppose the statement holds for some $i \in \{1, \dots, \Delta - 1\}$. That is, there are $T \leq k^{di(i-1)/2} b^{di}$ disjoint boxes $R_1 \cup \dots \cup I_R = [0, 1]^d$ such that $\mathcal{T}[i]$ is invariant over all μ within any given boxes (our inductive hypothesis). Fix one of these boxes R_t . We subdivide R_t into sub-boxes such that $\mathcal{T}[i+1]$ is invariant within each one of these smaller boxes. Let Q be any leaf of $\mathcal{T}[i]$, and for $\mu \in R_t$ let \mathcal{T}_μ denote the state of the tree using $\text{ascore}_{\mu_i}^i$ for each i at the point that Q is selected. Since $i < \Delta$, Q is not fathomed at line 4, regardless of μ . Next, since actions_i is path-wise for each i , the actions available at Q depend only on the path \mathcal{T}_Q from the root of \mathcal{T} to Q , which, by the inductive hypothesis, is invariant over all $\mu \in R_t$. That is, for all i $\text{actions}_i(\mathcal{T}_\mu, Q) = \text{actions}_i(\mathcal{T}_Q, Q)$ for all $\mu \in R_t$. Then, $\text{ascore}_{\mu_i}^i$ will select action $A_i \in \text{actions}_i(\mathcal{T}_Q, Q)$ if and only if

$$\begin{aligned}
A_i &= \operatorname{argmax}_{A_0 \in \text{actions}_i(\mathcal{T}_Q, Q)} \mu \cdot \text{ascore}_1^i(\mathcal{T}_\mu, Q, A_0) + (1 - \mu_i) \cdot \text{ascore}_2^i(\mathcal{T}_\mu, Q, A_0) \\
&= \operatorname{argmax}_{A_0 \in \text{actions}_i(\mathcal{T}_Q, Q)} \mu_i \cdot \text{ascore}_1^i(\mathcal{T}_Q, Q, A_0) + (1 - \mu_i) \cdot \text{ascore}_2^i(\mathcal{T}_Q, Q, A_0),
\end{aligned}$$

where the second equality follows from the fact that ascore_1^i and ascore_2^i are path-wise. Thus, for a fixed A_0 , $\text{ascore}_{\mu_i}^i$ is linear in μ_i , so for each A_0 there is at most one subinterval of $[0, 1]$ such that for all μ_i in that subinterval, A_0 maximizes $\text{ascore}_{\mu_i}^i$. Thus, each leaf of $\mathcal{T}[i]$ contributes at most b subintervals such that for μ_i within a given subinterval, the action of type i selected at each leaf of $\mathcal{T}[i]$ is invariant. $\mathcal{T}[i]$ consists of at most k^i leaves, so this is a total of at most $k^i b$ subintervals. Writing $R_t = I_1 \times \dots \times I_d$, we have established that for each i , there are $k^i b$ subintervals partitioning I_i into subintervals such that as μ_i varies over each subinterval, the action of type i selected at every leaf of $\mathcal{T}[i]$ is invariant. These subintervals partition R_t into at most $(k^i b)^d$ boxes. As before, since the actions selected at each leaf Q of $\mathcal{T}[i]$ are invariant, the set of children $\text{children}(\mathcal{T}_\mu, Q, A_1, \dots, A_d) = \text{children}(\mathcal{T}_Q, Q, A_1, \dots, A_d)$ of

Q added to the tree is also invariant, using the fact that `children` is path-wise. Therefore, within every sub-box of R_t , $\mathcal{T}[i+1]$ is invariant. The total number of boxes over each possible R_t is, by the induction hypothesis, at most $k^{di(i-1)/2}b^{di} \cdot k^{di}b^d = k^{d(i+1)i/2}b^{d(i+1)}$. ◀

The proof of Lemma 2 is identical in the multi-action setting. The proof of Lemma 4 is also identical: here, we fix a box R in the partition established in Lemma 9, and get an identical partition of $R \times [0, 1]$ such that the behavior of Algorithm 2 is invariant as λ varies in each subinterval of $[0, 1]$. The number of boxes in the final partition of $[0, 1]^{d+1}$ is $k^{d\Delta(\Delta-1)/2}b^{d\Delta} \cdot k^{5\Delta} \leq k^{d\Delta(9+\Delta)}b^{d\Delta}$. Our main pseudo-dimension bound for the multi-action setting follows from the same argument that exploits the framework of Balcan et al. [6].

► **Theorem 10.** *Let $\text{cost}(Q)$ be any tree-constant cost function, and let $\text{cost}_{\mu,\lambda}(Q)$ be the cost of the tree built by Algorithm 1 on input root node Q using action-selection scores parameterized by $\mu \in \mathbb{R}^d$, where $d = O(1)$, and node-selection score parameterized by λ . Then, $\text{Pdim}(\{\text{cost}_{\mu,\lambda}\}) = O(d\Delta^2 \log k + d\Delta \log b)$.*

When $d = O(1)$ we get the same pseudo-dimension bound as in the single-action setting: $\text{Pdim}(\{\text{cost}_{\mu,\lambda}\}) = O(\Delta^2 \log k + \Delta \log b)$, which is the statement of Theorem 6.