

Reinforcement Learning for Branch-and-Bound Optimisation Using Retrospective Trajectories

Christopher W. F. Parsonson^{1*}, Alexandre Laterre², Thomas D. Barrett²

¹UCL

²InstaDeep

Abstract

Combinatorial optimisation problems framed as mixed integer linear programmes (MILPs) are ubiquitous across a range of real-world applications. The canonical branch-and-bound algorithm seeks to exactly solve MILPs by constructing a search tree of increasingly constrained sub-problems. In practice, its solving time performance is dependent on heuristics, such as the choice of the next variable to constrain ('branching'). Recently, machine learning (ML) has emerged as a promising paradigm for branching. However, prior works have struggled to apply reinforcement learning (RL), citing sparse rewards, difficult exploration, and partial observability as significant challenges. Instead, leading ML methodologies resort to approximating high quality handcrafted heuristics with imitation learning (IL), which precludes the discovery of novel policies and requires expensive data labelling. In this work, we propose *retro branching*; a simple yet effective approach to RL for branching. By retrospectively deconstructing the search tree into multiple paths each contained within a sub-tree, we enable the agent to learn from shorter trajectories with more predictable next states. In experiments on four combinatorial tasks, our approach enables learning-to-branch without any expert guidance or pre-training. We outperform the current state-of-the-art RL branching algorithm by 3-5 \times and come within 20% of the best IL method's performance on MILPs with 500 constraints and 1000 variables, with ablations verifying that our retrospectively constructed trajectories are essential to achieving these results.

1 Introduction

A plethora of real-world problems fall under the broad category of combinatorial optimisation (CO) (vehicle routing and scheduling (Korte and Vygen 2012); protein folding (Perdomo-Ortiz et al. 2012); fundamental science (Barahona 1982)). Many CO problems can be formulated as mixed integer linear programmes (MILPs) whose task is to assign discrete values to a set of decision variables, subject to a mix of linear and integrality constraints, such that some objective function is maximised or minimised. The most popular method for finding exact solutions to MILPs is branch-and-bound (B&B) (Land and Doig 1960); a collection of heuris-

tics which increasingly tighten the bounds in which an optimal solution can reside (see Section 3). Among the most important of these heuristics is *variable selection* or *branching* (which variable to use to partition the chosen node's search space), which is key to determining B&B solve efficiency (Achterberg and Wunderling 2013).

State-of-the-art (SOTA) learning-to-branch approaches typically use the imitation learning (IL) paradigm to predict the action of a high quality but computationally expensive human-designed branching expert (Gasse et al. 2019). Since branching can be formulated as a Markov decision process (MDP) (He, Daumé, and Eisner 2014), reinforcement learning (RL) seems a natural approach. The long-term motivations of RL include the promise of learning novel policies from scratch without the need for expensive expert data, the potential to exceed expert performance without human design, and the capability to maximise the performance of a policy parameterised by an expressivity-constrained deep neural network (DNN).

However, branching has thus far proved largely intractable for RL for reasons we summarise into three key challenges. (1) *Long episodes*: Whilst even random branching policies are theoretically guaranteed to eventually find the optimal solution, poor decisions can result in episodes of tens of thousands of steps for the 500 constraint 1000 variable MILPs considered by Gasse et al. 2019. This raises the familiar RL challenges of reward sparsity (Trott et al. 2019), credit assignment (Harutyunyan et al. 2019), and high variance returns (Mao et al. 2019). (2) *Large state-action spaces*: Each branching step might have hundreds or thousands of potential branching candidates with a huge number of unique possible sub-MILP states. Efficient exploration to discover improved trajectories in such large state-action spaces is a well-known difficulty for RL (Agostinelli et al. 2019; Ecoffet et al. 2021). (3) *Partial observability*: When a branching decision is made, the next state given to the brancher is determined by the next sub-MILP visited by the node selection policy. Jumping around the B&B tree without the brancher's control whilst having only partial observability of the full tree makes the future states seen by the agent difficult to predict. Etheve et al. 2020 therefore postulated the benefit of keeping the MDP within a sub-tree to improve observability and introduced the SOTA fitting for minimising the sub-tree size (FMSTS) RL branching algorithm. However, in order to achieve this,

*Work undertaken during internship at InstaDeep.

Corresponding email: cwfparsonson@gmail.com
Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

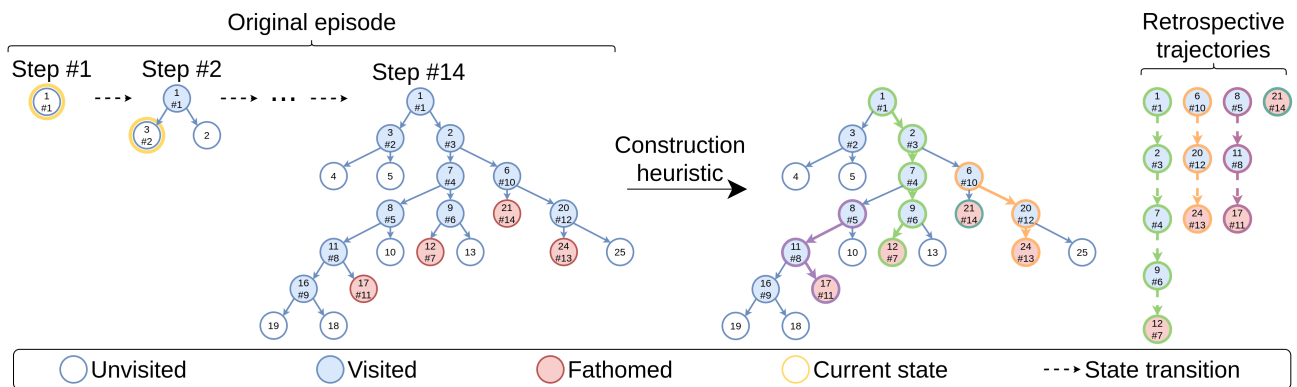


Figure 1: The proposed retro branching approach used during training. Each node is labelled with: Top: The unique ID assigned when it was added to the tree, and (where applicable); bottom: The step number (preceded by a ‘#’) at which it was visited by the brancher in the original Markov decision process (MDP). The MILP is first solved with the brancher and the B&B tree stored as usual (forming the ‘original episode’). Then, ignoring any nodes never visited by the agent, the nodes are added to trajectories using some ‘construction heuristic’ (see Sections 4 and 6) until each eligible node has been added to one, and only one, trajectory. Crucially, the order of the sequential states within a given trajectory may differ from the state visitation order of the original episode, but all states within the trajectory will be within the same sub-tree. These trajectories are then used for training.

FMSTS had to use a depth-first search (DFS) node selection policy which, as we demonstrate in Section 6, is highly sub-optimal and limits scalability.

In this work, we present *retro branching*; a simple yet effective method to overcome the above challenges and learn to branch via reinforcement. We follow the intuition of Etheve et al. (2020) that constraining each sequential MDP state to be within the same sub-tree will lead to improved observability. However, we posit that a branching policy taking the ‘best’ actions with respect to only the sub-tree in focus can still provide strong overall performance *regardless of the node selection policy used*. This is aligned with the observation that leading heuristics such as SB and PB also do not explicitly account for the node selection policy or predict how the global bound may change as a result of activity in other sub-trees. Assuming the validity of this hypothesis, we can discard the DFS node selection requirement of FMSTS whilst retaining the condition that sequential states seen during training must be within the same sub-tree.

Concretely, our retro branching approach (shown in Figure 1 and elaborated on in Section 4) is to, during training, take the search tree after the B&B instance has been solved and *retrospectively* select each subsequent state (node) to construct multiple trajectories. Each trajectory consists of sequential nodes within a single sub-tree, allowing the brancher to learn from shorter trajectories with lower return variance and more predictable future states. This approach directly addresses challenges (1) and (3) and, whilst the state-action space is still large, the shorter trajectories implicitly define more immediate auxiliary objectives relative to the tree. This reduces the difficulty of exploration since shorter trajectory returns will have a higher probability of being improved upon via stochastic action sampling than when a single long MDP is considered, thereby addressing (2). Furthermore, retro branching relieves the FMSTS requirement that the agent must be trained in a DFS node selection setting, en-

abling more sophisticated strategies to be used which are better suited for solving larger, more complex MILPs.

We evaluate our approach on MILPs with up to 500 constraints and 1000 variables, achieving a 3-5 \times improvement over FMSTS and coming within $\approx 20\%$ of the performance of the SOTA IL agent of Gasse et al. (2019). Furthermore, we demonstrate that, for small instances, retro branching can uncover policies superior to IL; a key motivation of using RL. Our results open the door to the discovery of new branching policies which can scale without the need for labelled data and which could, in principle, exceed the performance of SOTA handcrafted branching heuristics.

2 Related Work

Since the invention of B&B for exact CO by Land and Doig (1960), researchers have sought to design and improve the node selection (tree search), variable selection (branching), primal assignment, and pruning heuristics used by B&B, with comprehensive reviews provided by Achterberg 2007 and Tomazella and Nagano 2020. We focus on branching.

Classical branching heuristics. Pseudocost branching (PB) (Benichou et al. 1971) and strong branching (SB) (Aplegate et al. 1995, 2007) are two canonical branching algorithms. PB selects variables based on their historic branching success according to metrics such as bound improvement. Although the per-step decisions of PB are computationally fast, it must initialise the variable pseudocosts in some way which, if done poorly, can be particularly damaging to overall performance since early B&B decisions tend to be the most influential. SB, on the other hand, conducts a one-step lookahead for all branching candidates by computing their potential local dual bound gains before selecting the most favourable variable, and thus is able to make high quality decisions during the critical early stages of the search tree’s evolution. Despite its simplicity, SB is still today the best known policy for minimising the overall number of B&B

nodes needed to solve the problem instance (a popular B&B quality indicator). However, its computational cost renders SB infeasible in practice.

Learning-to-branch. Recent advances in deep learning have led machine learning (ML) researchers to contribute to exact CO (surveys provided by Lodi and Zarpellon 2017, Bengio, Lodi, and Prouvost 2021, and Cappart et al. 2021). Khalil et al. 2016 pioneered the community’s interest by using IL to train a support vector machine (SVM) to imitate the variable rankings of SB after the first 500 B&B node visits and thereafter use the SVM. Alvarez, Louveaux, and Wehenkel 2017 similarly imitated SB, but learned to predict the SB scores directly using Extremely Randomized Trees (Geurts, Ernst, and Wehenkel 2006). These approaches performed promisingly, but their per-instance training and use of SB at test time limited their scalability.

These issues were overcome by Gasse et al. 2019, who took as input a bipartite graph representation capturing the current B&B node state and predicted the corresponding action chosen by SB using a graph convolutional network (GCN). This alleviated the reliance on extensive feature engineering, avoided the use of SB at inference time, and demonstrated generalisation to larger instances than seen in training. Works since have sought to extend this method by introducing new observation features to generalise across heterogeneous CO instances (Zarpellon et al. 2021) and designing SB-on-a-GPU expert labelling methods for scalability (Nair et al. 2021).

Etheve et al. 2020 proposed FMSTS which, to the best of our knowledge, is the only published work to apply RL to branching and is therefore the SOTA RL branching algorithm. By using a DFS node selection strategy, they used the deep Q-network (DQN) approach (Mnih et al. 2013) to approximate the Q-function of the B&B sub-tree size rooted at the current node; a local Q-function which, in their setting, was equivalent to the number of global tree nodes. Although FMSTS alleviated issues with credit assignment and partial observability, it relied on using the DFS node selection policy (which can be far from optimal), was fundamentally limited by exponential sub-tree sizes produced by larger instances, and its associated models and data sets were not open-accessed.

3 Background

Mixed integer linear programming. An MILP is an optimisation task where values must be assigned to a set of n *decision variables* subject to a set of m *linear constraints* such that some linear *objective function* is minimised. MILPs can be written in the standard form

$$\arg \min_{\mathbf{x}} \{ \mathbf{c}^\top \mathbf{x} \mid \mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}, \mathbf{x} \in \mathbb{Z}^p \times \mathbb{R}^{n-p} \}, \quad (1)$$

where $\mathbf{c} \in \mathbb{R}^n$ is a vector of the objective function’s coefficients for each decision variable in \mathbf{x} such that $\mathbf{c}^\top \mathbf{x}$ is the objective value, $\mathbf{A} \in \mathbb{R}^{m \times n}$ is a matrix of the m constraints’ coefficients (rows) applied to n variables (columns), $\mathbf{b} \in \mathbb{R}^m$ is the vector of variable constraint right-hand side bound values which must be adhered to, and $\mathbf{l}, \mathbf{u} \in \mathbb{R}^n$ are the respective lower and upper variable value bounds. MILPs are

hard to solve owing to their *integrality constraint(s)* whereby $p \leq n$ decision variables must be an integer. If these integrality constraints are relaxed, the MILP becomes a linear programme (LP), which can be solved efficiently using algorithms such as simplex (Nelder and Mead 1965). The most popular approach for solving MILPs exactly is B&B.

Branch-and-bound. B&B is an algorithm composed of multiple heuristics for solving MILPs. It uses a search tree where nodes are MILPs and edges are partition conditions (added constraints) between them. Using a *divide and conquer* strategy, the MILP is iteratively partitioned into sub-MILPs with smaller solution spaces until an optimal solution (or, if terminated early, a solution with a worst-case optimality gap guarantee) is found. The task of B&B is to evolve the search tree until the provably optimal node is found.

Concretely, as summarised in Figure 2, at each step in the algorithm, B&B: (1) Selects an open (unfathomed leaf) node in the tree whose sub-tree seems promising to evolve; (2) selects (‘branches on’) a variable to tighten the bounds on the sub-MILP’s solution space by adding constraints either side of the variable’s LP solution value, generating two child nodes (sub-MILPs) beneath the focus node; (3) for each child, i) solve the relaxed LP (the *dual problem*) to get the *dual bound* (a bound on the best possible objective value in the node’s sub-tree) and, where appropriate, ii) solve the *primal problem* and find a feasible (but not necessarily optimal) solution satisfying the node’s constraints, thus giving the *primal bound* (the worst-case feasible objective value in the sub-tree); and (4) fathom any children (i.e. consider the sub-tree rooted at the child ‘fully known’ and therefore excluded from any further exploration) whose relaxed LP solution is integer-feasible, is worse than the incumbent (the globally best feasible node found so far), or which cannot meet the non-integrality constraints of the MILP. This process is repeated until the *primal-dual gap* (global primal-dual bound difference) is 0, at which point a provably optimal solution to the original MILP will have been found.

Note that the heuristics (i.e. primal, branching, and node selection) at each stage jointly determine the performance of B&B. More advanced procedures such as cutting planes (Mitchell 2009) and column generation (Barnhart et al. 1998) are available for enhancement, but are beyond the scope of this work. Note also that solvers such as SCIP 2022 only store ‘visitable’ nodes in memory, therefore in practice fathoming occurs at a feasible node where a branching decision led to the node’s two children being outside the established optimality bounds, being infeasible, or having an integer-feasible dual solution, thereby closing the said node’s sub-tree.

Q-learning. Q-learning is typically applied to sequential decision making problems formulated as an MDP defined by tuple $\{\mathcal{S}, \mathcal{U}, \mathcal{T}, \mathcal{R}, \gamma\}$. \mathcal{S} is a finite set of states, \mathcal{U} a set of actions, $\mathcal{T} : \mathcal{S} \times \mathcal{U} \times \mathcal{S} \rightarrow [0, 1]$ a transition function from state $s \in \mathcal{S}$ to $s' \in \mathcal{S}$ given action $u \in \mathcal{U}$, $\mathcal{R} : \mathcal{S} \rightarrow \mathbb{R}$ a function returning a scalar reward from s , and $\gamma \in [0, 1]$ a factor by which to discount expected future returns to their present value. It is an off-policy temporal difference method which aims to learn the action-value function mapping state-action pairs to the expected discounted sum of their immediate and future re-

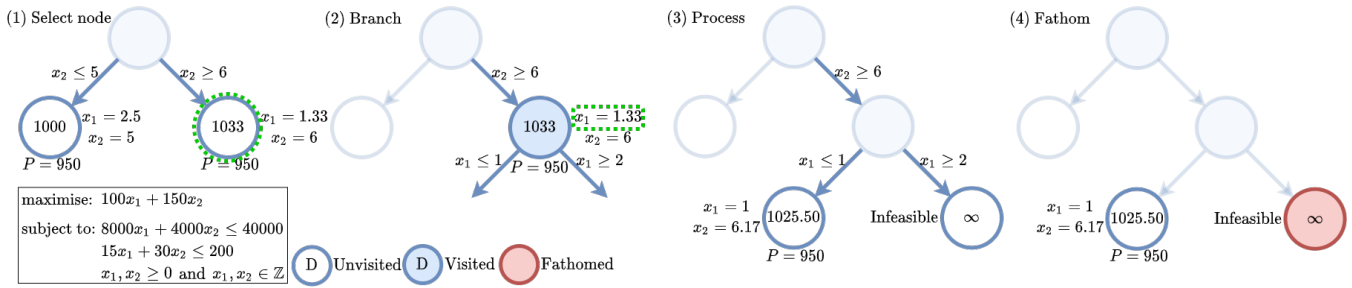


Figure 2: Typical 4-stage procedure iteratively repeated by B&B to solve an MILP. Each node represents an MILP derived from the original MILP being solved, and each edge represents the constraint added to derive a new child node (sub-MILP) from a given parent. Each node is labelled with the decision variable values of the solved LP relaxation on the right hand side, the corresponding dual bound in the centre, and the established primal bound beneath. Each edge is labelled with the introduced constraint to generate the child node. Green dotted outlines are used to indicate which node and variable were selected in stages (1) and (2) to lead to stages (3) and (4). The global primal (P) and dual (D) bounds are increasingly constrained by repeating stages 1-4 until P and D are equal, at which point a provably optimal solution will have been found. Note that for clarity we only show the detailed information needed at each stage, but that this does not indicate any change to the state of the tree.

wards when following a policy $\pi : \mathcal{S} \rightarrow \mathcal{U}$, $Q^\pi(s, u) = \mathbb{E}_\pi \left[\sum_{t'=t+1}^{\infty} \gamma^{t'-t} r(s_{t'}) | s_t = s, u_t = u \right]$. By definition, an optimal policy π_* will select an action which maximises the true Q-value $Q_*(s, u)$, $\pi_*(s) = \arg \max_{u'} Q_*(s, u')$. For scalability, DQN (Mnih et al. 2013; Hasselt, Guez, and Silver 2016) approximates this true Q-function using a DNN parameterised by θ such that $Q_\theta(s, u) \approx Q_*(s, u)$.

4 Retro Branching

We now describe our retro branching approach for learning-to-branch with RL.

States. At each time step t the B&B solver state is comprised of the search tree with past branching decisions, per-node LP solutions, the global incumbent, the currently focused leaf node, and any other solver statistics which might be tracked. To convert this information into a suitable input for the branching agent, we represent the MILP of the focus node chosen by the node selector as a bipartite graph. Concretely, the n variables and m constraints are connected by edges denoting which variables each constraint applies to. This formulation closely follows the approach of Gasse et al. 2019, with a full list of input features at each node detailed in Appendix E.

Actions. Given the MILP state s_t of the current focus node, the branching agent uses a policy $\pi(u_t | s_t)$ to select a variable u_t from among the p branching candidates.

Original full episode transitions. In the original full B&B episode, the next node visited is chosen by the node selection policy from amongst any of the open nodes in the tree. This is done independently of the brancher, which observes state information related only to the current focus node and the status of the global bounds. As such, the transitions of the ‘full episode’ are partially observable to the brancher, and it will therefore have the challenging task of needing to aggregate over unobservable states in external sub-trees to predict the long-term values of states and actions.

Retrospectively constructed trajectory transitions

(retro branching). To address the partial observability of the full episode, we retrospectively construct multiple trajectories where all sequential states in a given trajectory are within the same sub-tree, and where the trajectory’s terminal state is chosen from amongst the as yet unchosen fathomed sub-tree leaves. A visualisation of our approach is shown in Figure 1. Concretely, during training, we first solve the instance as usual with the RL brancher and any node selection heuristic to form the ‘original episode’. When the instance is solved, rather than simply adding the originally observed MDP’s transitions to the DQN replay buffer, we retrospectively construct multiple trajectory paths through the search tree. This construction process is done by starting at the highest level node not yet added to a trajectory, selecting an as yet unselected fathomed leaf in the sub-tree rooted at said node using some ‘construction heuristic’ (see Section 6), and using this root-leaf pair as a source-destination with which to construct a path (a ‘retrospective trajectory’). This process is iteratively repeated until each eligible node in the original search tree has been added to one, and only one, retrospective trajectory. The transitions of each trajectory are then added to the experience replay buffer for learning. Note that retrospective trajectories are only used during training, therefore retro branching agents have no additional inference-time overhead.

Crucially, retro branching determines the sequence of states in each trajectory (i.e. the transition function of the MDP) such that the next state(s) observed in a given trajectory will *always* be within the same sub-tree (see Figure 1) regardless of the node selection policy used in the original B&B episode. Our reasoning behind this idea is that the state(s) beneath the current focus node within its sub-tree will have characteristics (bounds, introduced constraints, etc.) which are strongly related with those of the current node, making them more observable than were the next states to be chosen from elsewhere in the search tree, as can occur in the ‘original B&B’ episode. Moreover, by correlating the agent’s maximum trajectory length with the depth of the tree rather than the total number of nodes, reconstructed trajectories

have orders of magnitude fewer steps and lower return variance than the original full episode, making learning tractable on large MILPs. Furthermore, because the sequential nodes visited are chosen retrospectively in each trajectory, unlike with FMSTS, any node selection policy can be used during training. As we show in Section 6, this is a significant help when solving large and complex MILPs.

Rewards. As demonstrated in Section 6, the use of reconstructed trajectories enables a simple distance-to-goal reward function to be used; a $r = -1$ punishment is issued to the agent at each step except when the agent’s action fathomed the sub-tree, where the agent receives $r = 0$. This provides an incentive for the the branching agent to reach the terminal state as quickly as possible. When aggregated over all trajectories in a given sub-tree, this auxiliary objective corresponds to fathoming the whole sub-tree (and, by extension, solving the MILP) in as few steps as possible. This is because the only nodes which are stored by SCIP 2022 and which the brancher will be presented with will be feasible nodes which *potentially* contain the optimal solution beneath them. As such, any action chosen by the brancher which provably shows either the optimal solution to not be beneath the current node or which finds an integer feasible dual solution (i.e. an action which fathoms the sub-tree beneath the node) will be beneficial, because it will prevent SCIP from being able to further needlessly explore the node’s sub-tree.

A note on partial observability. In the above retrospective formulation of the branching MDP, the primal, branching, and node selection heuristics active in other sub-trees will still influence the future states and fathoming conditions of a given retrospective trajectory. We posit that there are two extremes; DFS node selection where future states are fully observable to the brancher, and non-DFS node selection where they are heavily obscured. As shown in Section 6, our retrospective node selection setting strikes a balance between these two extremes, attaining sufficient observability to facilitate learning while enabling the benefits of short, low variance trajectories with sophisticated node selection strategies which make handling larger MILPs tractable.

5 Experimental Setup

All code for reproducing the experiments and links to the generated data sets are provided at https://github.com/cwfpinson/retro_branching.

Network architecture and learning algorithm. We used the GCN architecture of Gasse et al. 2019 to parameterise the DQN value function with some minor modifications which we found to be helpful (see Appendix B.1). We trained our network with n-step DQN (Sutton 1988; Mnih et al. 2013) using prioritised experience replay (Schaul et al. 2016), soft target network updates (Lillicrap et al. 2016), and an epsilon-stochastic exploration policy (see Appendix A.1 for a detailed description of our RL approach and the corresponding algorithms and hyperparameters used).

B&B environment. We used the open-source Ecole (Prouvost et al. 2020) and PySCIPOpt (Maher et al. 2016) libraries with SCIP 7.0.1 (SCIP 2022) as the backend solver to do instance generation and testing. Where possible, we used the training and testing protocols of Gasse et al. (2019).

MILP Problem classes. In total, we considered four NP-hard problem benchmarks: set covering (Balas, Ho, and Center 2018), combinatorial auction (Leyton-Brown, Pearson, and Shoham 2000), capacitated facility location (Litvinchev and Ozuna Espinosa 2012), and maximum independent set (Bergman et al. 2016).

Baselines. We compared retro branching against the SOTA FMSTS RL algorithm of Etheve et al. (2020) (see Appendix F for implementation details) and the SOTA IL approach of Gasse et al. (2019) trained and validated with 100 000 and 20 000 strong branching samples respectively. For completeness, we also compared against the SB heuristic imitated by the IL agent, the canonical PB heuristic, and a random brancher (equivalent in performance to most infeasible branching (Achterberg, Koch, and Martin 2004)). Note that we have omitted direct comparison to the SOTA tuned commercial solvers, which we do not claim to be competitive with at this stage. To evaluate the quality of the agents’ branching decisions, we used 100 validation instances (see Appendix C for an analysis of this data set size) which were unseen during training, reporting the total number of tree nodes and LP iterations as key metrics to be minimised.

6 Results & Discussion

6.1 Performance of Retro Branching

Comparison to the SOTA RL branching heuristics. We considered set covering instances with 500 rows and 1000 columns. To demonstrate the benefit of the proposed retro branching method, we trained a baseline ‘Original’ agent on the original full episode, receiving the same reward as our retro branching agent (-1 at each non-terminal step and 0 for a terminal action which ended the episode). We also trained the SOTA RL FMSTS branching agent in a DFS setting and, at test time, validated the agent in both a DFS (‘FMSTS-DFS’) and non-DFS (‘FMSTS’) environment to fairly compare the policies. Note that the FMSTS agent serves as an ablation to analyse the influence of training on retrospective trajectories, since it uses our auxiliary objective but without retrospective trajectories, and that the Original agent further ablates the auxiliary objective since its ‘terminal step’ is defined as ending the B&B episode (where it receives $r_t = 0$ rather than $r_t = -1$). As shown in Figure 3a, the Original agent was unable to learn on these large instances, with retro branching achieving $14\times$ fewer nodes at test time. FMSTS also performed poorly, with highly unstable learning and a final performance $5\times$ and $3\times$ poorer than retro branching in the DFS and non-DFS settings respectively (see Figure 3c). We posit that the cause of the poor FMSTS performance is due to its use of the sub-optimal DFS node selection policy, which is ill-suited for handling large MILPs and results in $\approx 10\%$ of episodes seen during training being on the order of 10-100k steps long (see Figure 3b), which makes learning significantly harder for RL.

Comparison to non-RL branching heuristics. Having demonstrated that the proposed retro branching method makes learning-to-branch at scale tractable for RL, we now compare retro branching with the baseline branchers to understand the efficacy of RL in the context of the current literature.

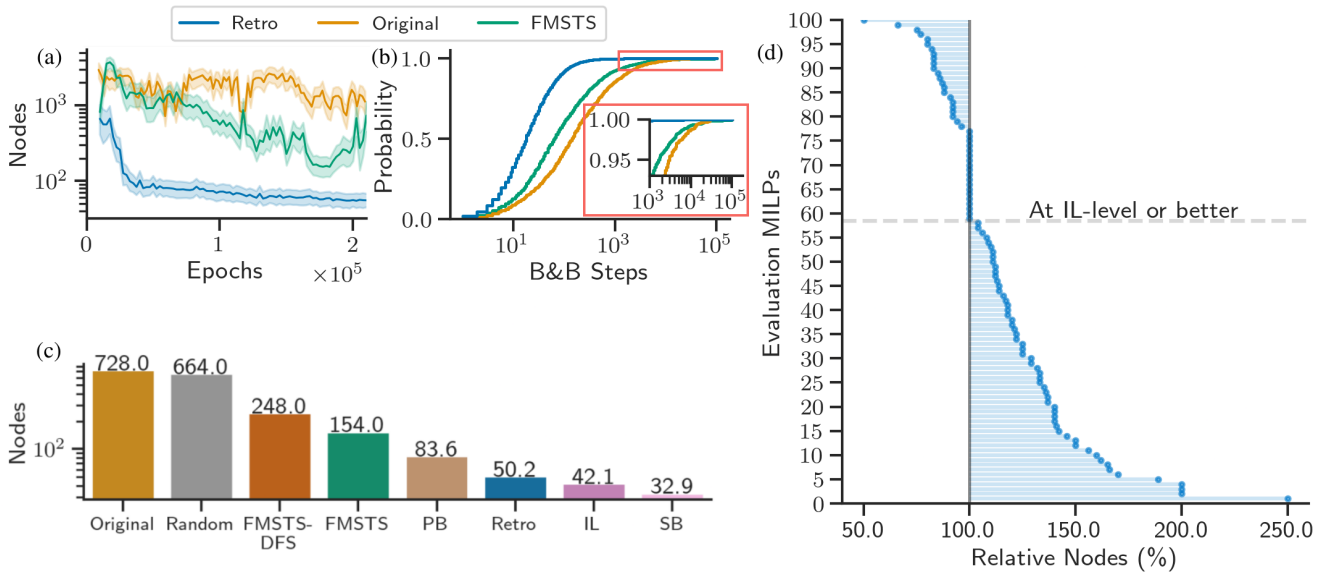


Figure 3: Performances of the branching agents on the 500×1000 set covering instances. (a) Validation curves for the RL agents evaluated in the same non-DFS setting. (b) CDF of the number of B&B steps taken by the RL agents for each instance seen during training. (c) The best validation performances of each branching agent. (d) The instance-level validation performance of the retro branching agent relative to the IL agent, with RL matching or beating IL on 42% of test instances.

Figure 3c shows how retro branching compares to other policies on large 500×1000 set covering instances. While the agent outperforms PB, it only matches or beats IL on 42% of the test instances (see Figure 3d) and, on average, has a $\approx 20\%$ larger B&B tree size. Therefore although our RL agent was still improving and was limited by compute (see Appendix A.2), and in spite of our method outperforming the current SOTA FMSTS RL brancher, RL has not yet been able to match or surpass the SOTA IL agent at scale. This will be an interesting area of future work, as discussed in Section 7.

6.2 Analysis of Retro Branching

Verifying that RL can outperform IL. In addition to not needing labelled data, a key motivation for using RL over IL for learning-to-branch is the potential to discover superior policies. While Figure 3 showed that, at test-time, retro branching matched or outperformed IL on 42% of instances, IL still had a lower average tree size. As shown in Table 1, we found that, on small set covering instances with 165 constraints and 230 variables, RL could outperform IL by $\approx 20\%$. While improvement on problems of this scale is not the primary challenge facing ML-B&B solvers, we are encouraged by this demonstration that it is possible for an RL agent to learn a policy better able to maximise the performance of an expressivity-constrained network than imitating an expert such as SB without the need for pre-training or expensive data labelling procedures (see Appendix H).

For completeness, Table 1 also compares the retro branching agent to the IL, PB, and SB branching policies evaluated on 100 unseen instances of four NP-hard CO benchmarks. We considered instances with 10 items and 50 bids for combinatorial auction, 5 customers and facilities for capacitated

facility location, and 25 nodes for maximum independent set. RL achieved a lower number of tree nodes than PB and IL on all problems except combinatorial auction. This highlights the potential for RL to learn improved branching policies to solve a variety of MILPs.

Demonstrating the independence of retro branching to future state selection. As described in Section 4, in order to retrospectively construct a path through the search tree, a fathomed leaf node must be selected. We refer to the method for selecting the leaf node as the *construction heuristic*. The future states seen by the agent are therefore determined by the construction heuristic (used in training) and the node selection heuristic (used in training and inference).

During our experiments, we found that the specific construction heuristic used had little impact on the performance of our agent. Figure 4a shows the validation curves for four agents trained on 500×1000 set covering instances each using one of the following construction heuristics: Maximum LP gain ('MLPG': Select the leaf with the largest LP gain); random ('R': Randomly select a leaf); visitation order ('VO': Select the leaf which was visited first in the original episode); and deepest ('D': Select the leaf which results in the longest trajectory). As shown, all construction heuristics resulted in roughly the same performance (with MLPG performing only slightly better). This suggests that the agent learns to reduce the trajectory length regardless of the path chosen by the construction heuristic. Since the specific path chosen is independent of node selection, we posit that the relative strength of an RL agent trained with retro branching will also be independent of the node selection policy used.

To test this, we took our best retro branching agent trained with the default SCIP node selection heuristic and tested it

| Method | Set Covering | | Combinatorial Auction | | Capacitated Facility Location | | Maximum Independent Set | |
|--------|--------------|-------------|-----------------------|-------------|-------------------------------|-------------|-------------------------|-------------|
| | # LPs | # Nodes | # LPs | # Nodes | # LPs | # Nodes | # LPs | # Nodes |
| SB | 184 | 6.76 | 13.2 | 4.64 | 28.2 | 10.2 | 19.2 | 3.80 |
| PB | 258 | 12.8 | 22.0 | 7.80 | 28.0 | 10.2 | 25.4 | 5.77 |
| IL | 244 | 10.5 | 16.0 | 5.29 | 28.0 | 10.2 | 20.1 | 4.08 |
| Retro | 206 | 8.68 | 18.1 | 5.73 | 28.4 | 10.1 | 19.1 | 4.01 |

Table 1: Test-time comparison of the best agents on the evaluation instances of the four NP-hard small CO problems considered.




Figure 4: 500×1000 set covering performances. (a) Validation curves for four retro branching agents each trained with a different trajectory construction heuristic: Maximum LP gain (MLPG); random (R); visitation order (VO); and deepest (D). (b) The performances of the best retro branching agent deployed in three different node selection environments (default SCIP, DFS, and breadth-first search (BFS)) normalised relative to the performances of PB (measured by number of tree nodes).

on the 500×1000 validation instances in the default, DFS, and breadth-first search (BFS) SCIP node selection settings. To make the performances of the brancher comparable across these settings, we normalised the mean tree sizes with those of PB (a branching heuristic independent of the node selector) to get the performance relative to PB in each environment. As shown in Figure 4b, our agent achieved consistent relative performance regardless of the node selection policy used, indicating its indifference to the node selector.

7 Conclusions

We have introduced retro branching; a retrospective approach to constructing B&B trajectories in order to aid learning-to-branch with RL. We posited that retrospective trajectories address the challenges of long episodes, large state-action spaces, and partially observable future states which otherwise make branching an acutely difficult task for RL. We empirically demonstrated that retro branching outperforms the current SOTA RL method by $3\text{-}5\times$ and comes within 20% of the performance of IL whilst matching or beating it on 42% of test instances. Moreover, we showed that RL can surpass the performance of IL on small instances, exemplifying a key advantage of RL in being able to discover novel performance-maximising policies for expressivity-constrained networks without the need for pre-training or expert examples. However, retro branching was not able to exceed the IL agent at scale. In this section we outline areas of further work.

Partial observability. A limitation of our proposed approach is the remaining partial observability of the MDP, with activity external to the current sub-tree and branching

decision influencing future bounds, states, and rewards. In this and other studies, variable and node selection have been considered in isolation. An interesting approach would be to combine node and variable selection, giving the agent full control over how the B&B tree is evolved.

Reward function. The proposed trajectory reconstruction approach can facilitate a simple RL reward function which would otherwise fail were the original ‘full’ tree episode used. However, assigning a -1 reward at each step in a given trajectory ignores the fact that certain actions, particularly early on in the B&B process, can have significant influence over the length of multiple trajectories. This could be accounted for in the reward signal, perhaps by using a retrospective backpropagation method (similar to value backpropagation in Monte Carlo tree search (Silver et al. 2016, 2017)).

Exploration. The large state-action space and the complexity of making thousands of sequential decisions which together influence final performance in complex ways makes exploration in B&B an acute challenge for RL. One reason for RL struggling to close the 20% performance gap with IL at scale could be that, at some point, stochastic action sampling to explore new policies is highly unlikely to find trajectories with improved performance. As such, more sophisticated exploration strategies could be promising, such as novel experience intrinsic reward signals (Burda et al. 2019; Zhang et al. 2021), reverse backtracking through the episode to improve trajectory quality (Salimans and Chen 2018; Agostinelli et al. 2019; Ecoffet et al. 2021), and avoiding local optima using auxiliary distance-to-goal rewards (Trott et al. 2019) or evolutionary strategies (Conti et al. 2018).

Acknowledgments

We would like to thank Maxime Gasse, Antoine Prouvost, and the rest of the Ecole development team for answering our questions on SCIP and Ecole, and also the anonymous reviewers for their constructive comments on earlier versions of this paper.

References

- Achterberg, T. 2007. *Constraint Integer Programming*. Doctoral thesis, Technische Universität Berlin, Fakultät II - Mathematik und Naturwissenschaften, Berlin.
- Achterberg, T.; Koch, T.; and Martin, A. 2004. Branching rules revisited. Technical Report 04-13, ZIB, Takustr. 7, 14195 Berlin.
- Achterberg, T.; and Wunderling, R. 2013. *Mixed Integer Programming: Analyzing 12 Years of Progress*. Facets of Combinatorial Optimization.
- Agostinelli, F.; McAleer, S.; Shmakov, A.; and Baldi, P. 2019. Solving the Rubik's cube with deep reinforcement learning and search. *Nature Machine Intelligence*, 1(8): 356–363.
- Alvarez, A. M.; Louveaux, Q.; and Wehenkel, L. 2017. A Machine Learning-Based Approximation of Strong Branching. *INFORMS J. Comput.*, 29: 185–195.
- Applegate, D. L.; Bixby, R. E.; Chvatal, V.; and Cook, W. J. 2007. *The Traveling Salesman Problem: A Computational Study*. Princeton University Press.
- Applegate, D. L.; Bixby, R. E.; Chvatal, V.; and Cook, W. J. 1995. Finding cuts in the TSP (A preliminary report). Technical report, DIMACS.
- Balas, E.; Ho, A.; and Center, C. M. U. R. 2018. Set covering algorithms using cutting planes, heuristics, and subgradient optimization : a computational study. *Defense Technical Information Center*.
- Barahona, F. 1982. On the computational complexity of Ising spin glass models. *Journal of Physics A: Mathematical and General*, 15(10): 3241.
- Barnhart, C.; Johnson, E. L.; Nemhauser, G. L.; Savelsbergh, M. W. P.; and Vance, P. H. 1998. Branch-And-Price: Column Generation for Solving Huge Integer Programs. *Oper. Res.*, 46(3): 316–329.
- Bengio, Y.; Lodi, A.; and Prouvost, A. 2021. Machine learning for combinatorial optimization: A methodological tour d'horizon. *European Journal of Operational Research*, 290(2): 405–421.
- Benichou, M.; Gauthier, J. M.; Girodet, P.; Hentges, G.; Ribiere, G.; and Vincent, O. 1971. Experiments in mixed-integer linear programming. *Mathematical Programming*, 1(1): 76–94.
- Bergman, D.; Cire, A. A.; Hoeve, W.-J. v.; and Hooker, J. 2016. *Decision Diagrams for Optimization*. Springer Publishing Company, Incorporated, 1st edition. ISBN 3319428470.
- Burda, Y.; Edwards, H.; Storkey, A.; and Klimov, O. 2019. Exploration by random network distillation. In *International Conference on Learning Representations*.
- Cappart, Q.; Chételat, D.; Khalil, E. B.; Lodi, A.; Morris, C.; and Veličković, P. 2021. Combinatorial Optimization and Reasoning with Graph Neural Networks. In Zhou, Z.-H., ed., *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, 4348–4355. International Joint Conferences on Artificial Intelligence Organization. Survey Track.
- Conti, E.; Madhavan, V.; Such, F. P.; Lehman, J.; Stanley, K. O.; and Clune, J. 2018. Improving Exploration in Evolution Strategies for Deep Reinforcement Learning via a Population of Novelty-Seeking Agents. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS'18*, 5032–5043. Red Hook, NY, USA: Curran Associates Inc.
- Ecoffet, A.; Huizinga, J.; Lehman, J.; Stanley, K. O.; and Clune, J. 2021. First return, then explore. *Nature*, 590(7847): 580–586.
- Etheve, M.; Alès, Z.; Bissuel, C.; Juan, O.; and Kedad-Sidhoum, S. 2020. Reinforcement Learning for Variable Selection in a Branch and Bound Algorithm. *Lecture Notes in Computer Science*, 176–185.
- Gasse, M.; Chételat, D.; Ferroni, N.; Charlin, L.; and Lodi, A. 2019. *Exact Combinatorial Optimization with Graph Convolutional Neural Networks*. Red Hook, NY, USA: Curran Associates Inc.
- Geurts, P.; Ernst, D.; and Wehenkel, L. 2006. Extremely randomized trees. *Machine Learning*, 63(1): 3–42.
- Harutyunyan, A.; Dabney, W.; Mesnard, T.; Heess, N.; Azar, M. G.; Piot, B.; van Hasselt, H.; Singh, S.; Wayne, G.; Precup, D.; and Munos, R. 2019. *Hindsight Credit Assignment*. Red Hook, NY, USA: Curran Associates Inc.
- Hasselt, H. v.; Guez, A.; and Silver, D. 2016. Deep Reinforcement Learning with Double Q-Learning. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, AAAI'16*, 2094–2100. AAAI Press.
- He, H.; Daumé, H.; and Eisner, J. 2014. Learning to Search in Branch-and-Bound Algorithms. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS'14*, 3293–3301. Cambridge, MA, USA: MIT Press.
- Khalil, E. B.; Bodic, P. L.; Song, L.; Nemhauser, G. L.; and Dilkina, B. N. 2016. Learning to Branch in Mixed Integer Programming. In *AAAI*.
- Korte, B. H.; and Vygen, J. 2012. *Combinatorial Optimization: Theory and Algorithms*. New York, NY: Springer-Verlag. ISBN 9783642244889 3642244882 3642244874 9783642244872.
- Land, A. H.; and Doig, A. G. 1960. An Automatic Method of Solving Discrete Programming Problems. *Econometrica*, 28(3): pp. 497–520.
- Leyton-Brown, K.; Pearson, M.; and Shoham, Y. 2000. Towards a Universal Test Suite for Combinatorial Auction Algorithms. In *Proceedings of the 2nd ACM Conference on Electronic Commerce, EC '00*, 66–76. New York, NY, USA: Association for Computing Machinery. ISBN 1581132727.

- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2016. Continuous control with deep reinforcement learning. In Bengio, Y.; and LeCun, Y., eds., *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Litvinchev, I.; and Ozuna Espinosa, E. L. 2012. Solving the Two-Stage Capacitated Facility Location Problem by the Lagrangian Heuristic. In Hu, H.; Shi, X.; Stahlbock, R.; and Voß, S., eds., *Computational Logistics*, 92–103. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-642-33587-7.
- Lodi, A.; and Zarpellon, G. 2017. On learning and branching: a survey. *TOP*, 25(2): 207–236.
- Maher, S.; Miltenberger, M.; Pedroso, J. P.; Rehfeldt, D.; Schwarz, R.; and Serrano, F. 2016. PySCIPOpt: Mathematical Programming in Python with the SCIP Optimization Suite. In *Mathematical Software – ICMS 2016*, 301–307. Springer International Publishing.
- Mao, H.; Venkatakrishnan, S. B.; Schwarzkopf, M.; and Alizadeh, M. 2019. Variance Reduction for Reinforcement Learning in Input-Driven Environments. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Mitchell, J. E. 2009. *Integer programming: branch and cut algorithms*. *Integer Programming: Branch and Cut Algorithms*, 1643–1650. Boston, MA: Springer US. ISBN 978-0-387-74759-0.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. 2013. Playing Atari with Deep Reinforcement Learning. *NIPS Deep Learning Workshop*.
- Nair, V.; Bartunov, S.; Gimeno, F.; von Glehn, I.; Lichocki, P.; Lobov, I.; O’Donoghue, B.; Sonnerat, N.; Tjandraatmadja, C.; Wang, P.; Addanki, R.; Hapuarachchi, T.; Keck, T.; Keeling, J.; Kohli, P.; Ktena, I.; Li, Y.; Vinyals, O.; and Zwols, Y. 2021. Solving Mixed Integer Programs Using Neural Networks. *arXiv*.
- Nelder, J. A.; and Mead, R. 1965. A simplex method for function minimization. *Computer Journal*, 7: 308–313.
- Perdomo-Ortiz, A.; Dickson, N.; Drew-Brook, M.; Rose, G.; and Aspuru-Guzik, A. 2012. Finding low-energy conformations of lattice protein models by quantum annealing. *Scientific Reports*, 2: 571.
- Prouvost, A.; Dumouchelle, J.; Scavuzzo, L.; Gasse, M.; Chételat, D.; and Lodi, A. 2020. Ecole: A Gym-like Library for Machine Learning in Combinatorial Optimization Solvers. In *Learning Meets Combinatorial Algorithms at NeurIPS2020*.
- Salimans, T.; and Chen, R. 2018. Learning Montezuma’s Revenge from a Single Demonstration. *NeurIPS Deep Reinforcement Learning Workshop*.
- Schaul, T.; Quan, J.; Antonoglou, I.; and Silver, D. 2016. Prioritized Experience Replay. In *International Conference on Learning Representations*.
- SCIP. 2022. The SCIP Optimization Suite 7.0. *Optimization Online*.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; van den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; Dieleman, S.; Grewe, D.; Nham, J.; Kalchbrenner, N.; Sutskever, I.; Lillicrap, T.; Leach, M.; Kavukcuoglu, K.; Graepel, T.; and Hassabis, D. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587): 484–489.
- Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; Chen, Y.; Lillicrap, T.; Hui, F.; Sifre, L.; van den Driessche, G.; Graepel, T.; and Hassabis, D. 2017. Mastering the game of Go without human knowledge. *Nature*, 550(7676): 354–359.
- Sutton, R. S. 1988. Learning to predict by the methods of temporal differences. *Machine Learning*, 3(1): 9–44.
- Tomazella, C. P.; and Nagano, M. S. 2020. A comprehensive review of Branch-and-Bound algorithms: Guidelines and directions for further research on the flowshop scheduling problem. *Expert Systems with Applications*, 158: 113556.
- Trott, A.; Zheng, S.; Xiong, C.; and Socher, R. 2019. Keeping Your Distance: Solving Sparse Reward Tasks Using Self-Balancing Shaped Rewards. In *NeurIPS*.
- Zarpellon, G.; Jo, J.; Lodi, A.; and Bengio, Y. 2021. Parameterizing Branch-and-Bound Search Trees to Learn Branching Policies. *arXiv:2002.05120*.
- Zhang, T.; Xu, H.; Wang, X.; Wu, Y.; Keutzer, K.; Gonzalez, J. E.; and Tian, Y. 2021. NovelD: A Simple yet Effective Exploration Criterion. In Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*.