






Local Branching Relaxation Heuristics for Integer Linear Programs

Taoan Huang¹(✉) , Aaron Ferber¹ , Yuandong Tian² , Bistra Dilkina¹ ,
and Benoit Steiner² 

¹ University of Southern California, Los Angeles, USA

{taoanhua, aferber, dilkina}@usc.edu

² Meta AI (FAIR), Menlo Park, USA

yuandong@meta.com

Abstract. Large Neighborhood Search (LNS) is a popular heuristic algorithm for solving combinatorial optimization problems (COP). It starts with an initial solution to the problem and iteratively improves it by searching a large neighborhood around the current best solution. LNS relies on heuristics to select neighborhoods to search in. In this paper, we focus on designing effective and efficient heuristics in LNS for integer linear programs (ILP) since a wide range of COPs can be represented as ILPs. Local Branching (LB) is a heuristic that selects the neighborhood that leads to the largest improvement over the current solution in each iteration of LNS. LB is often slow since it needs to solve an ILP of the same size as input. Our proposed heuristics, LB-RELAX and its variants, use the linear programming relaxation of LB to select neighborhoods. Empirically, LB-RELAX and its variants compute as effective neighborhoods as LB but run faster. They achieve state-of-the-art any-time performance on several ILP benchmarks.

Keywords: Integer Linear Program · Large Neighborhood Search · Heuristic Search

1 Introduction

Combinatorial optimization problems (COP) concerns a wide variety of real-world applications, including vehicle routing [42], path planning [35] and resource allocation [34] problems. Many of them are difficult to solve with limited computational resources due to their NP-Hardness. Nonetheless, the widespread importance of COPs has inspired research in designing algorithms for solving them, including exact algorithms, approximation algorithms, heuristic algorithms and data-driven algorithms.

In this paper, we focus specifically on Integer Linear Programs (ILPs) since it is a powerful tool to model and solve a broad collection of COPs, including graph optimization [40], mechanism design [11], facility location [4, 19] and network design [12, 21] problems. Branch-and-Bound (BnB) is an optimal and complete

tree search algorithm and is one of the state-of-the-art algorithms for ILPs [27]. It is also the core of many ILP solvers such as SCIP [8] and Gurobi [17]. Huge research effort has been made to improve it over the past decades [2]. However, BnB still falls short of delivering practical impact due to scalability issues [14, 24]. On the other hand, Large Neighborhood Search (LNS) is a powerful heuristic algorithm for hard COPs and has been recently applied to solve ILPs [40, 41, 43] in the machine learning (ML) community.

To solve ILPs, LNS starts with an initial solution, i.e., a feasible assignment of values to the variables. It then iteratively improves the best solution found so far (i.e., the *incumbent solution*), by applying *destroy heuristics* to select a subset of variables and solving a sub-ILP that optimizes only the selected variables while leaving others fixed. ML-based destroy heuristics are shown to be efficient and effective but they are often tailored for a specific problem domain and require extensive computational resources for learning. A few non-ML destroy heuristics have been studied, such as the randomized heuristics [40, 41] and the Local Branching (LB) heuristic [13, 41], but they are either less efficient or effective compared to the ML-based ones. The randomized heuristics select the neighborhood by quickly randomly sampling a subset of variables which is often of bad quality. LB computes the optimal solution across all possible search neighborhoods that differs from the current incumbent solutions on a limited number of variables; however, LB is computationally expensive since it requires solving an ILP that has the same size as the original problem.

To strike a balance between efficiency and effectiveness, we propose a simple yet effective destroy heuristic LB-RELAX that is based on the linear programming (LP) relaxation of LB. Instead of solving an ILP to find the neighborhood as LB does, LB-RELAX computes its LP relaxation. It then selects the variables greedily based on the difference between the values in the incumbent solution and the LP relaxation solution. We also propose two other variants, LB-RELAX-S and LB-RELAX-R, that deploy a sampling method and combine the randomized heuristic with LB-RELAX to help escape local optima more efficiently, respectively. In experiments, we compare LB-RELAX and its variants against LNS with baseline destroy heuristics and BnB on several ILP benchmarks and show that they achieve state-of-the-art anytime performance. We also show that LB-RELAX achieves competitive results with, sometimes even outperform, the ML-based destroy heuristics. We also test LB-RELAX and its variants on selected difficult MIPLIB instances [16] that encompass diverse problem domains, structures and sizes and show that they achieve best performance on at least 40% of the instances. We also empirically show that LB-RELAX and LB-RELAX-S find neighborhoods of similar quality but is much faster than LB. They sometimes even outperform LB due to LB being too slow to find good enough neighborhoods within a reasonable time cutoff.

2 Background

In this section, we first define ILP and introduce its LP relaxation. We then introduce LNS for ILP solving and the Local Branching (LB) heuristic.

Algorithm 1. LNS for ILPs

```

1: Input: An ILP.
2:  $\mathbf{x}^0 \leftarrow$  Find an initial solution to the input ILP
3:  $t \leftarrow 0$ 
4: while time limit not exceeded do
5:    $\mathcal{X}^t \leftarrow$  Select a subset of variables to destroy
6:    $\mathbf{x}^{t+1} \leftarrow$  Solve the ILP with additional constraints  $\{x_i = x_i^t : x_i \notin \mathcal{X}^t\}$ 
7:    $t \leftarrow t + 1$ 
8: return  $\mathbf{x}^t$ 

```

2.1 ILP and Its LP Relaxation

An *integer linear program (ILP)* is defined as

$$\min \mathbf{c}^\top \mathbf{x} \quad \text{s.t. } \mathbf{A}\mathbf{x} \leq \mathbf{b} \text{ and } \mathbf{x} \in \{0, 1\}^n,$$

where $\mathbf{x} = (x_1, \dots, x_n)^\top$ denotes the n binary variables to be optimized, $\mathbf{c} \in \mathbb{R}^n$ denotes the vector of objective coefficients and $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$ specify m linear constraints. A *solution* to the ILP is an feasible assignment of values to the variables.

The *linear programming (LP) relaxation* of an ILP is obtained by relaxing binary variables in the ILP to continuous variables between 0 and 1, i.e., by replacing the integer constraint $\mathbf{x} \in \{0, 1\}^n$ with $\mathbf{x} \in [0, 1]^n$.

Note that, in this paper, we focus on the formulation above that consists of only binary variables, but our methods can also be applied to mixed integer linear programs with continuous variables and/or non-binary integer variables.

2.2 LNS for ILP Solving

LNS is a heuristic algorithm that starts with an initial solution and then iteratively reoptimizes a part of the solution by applying the destroy and repair operations until a time limit is exceeded. Let \mathbf{x}^0 be the initial solution. In iteration $t \geq 0$ of the LNS, given the *incumbent solution* \mathbf{x}^t , defined as the best solution found so far, a destroy operation is done by a *destroy heuristic* where it selects a subset of k_t variables $\mathcal{X}^t = \{x_{i_1}, \dots, x_{i_{k_t}}\}$. The repair operation is done by solving a sub-ILP with \mathcal{X}^t being the variables while fixing the values of $x_j \notin \mathcal{X}^t$ to be the same as in \mathbf{x}^t . Compared to BnB, LNS is more effective in improving the objective value $\mathbf{c}^\top \mathbf{x}$, or the primal bound, especially on difficult instances [40, 41, 43]. Compared to other local search methods, LNS explores a large neighborhood in each step and thus, is more effective in avoiding local minima. LNS for ILPs is summarized in Algorithm 1.

2.3 LB Heuristic

The LB Heuristic [13] is originally proposed as a primal heuristic in BnB but is also applicable in LNS for ILP solving [31, 41]. Given the incumbent solution \mathbf{x}^t

in iteration t of LNS, the LB heuristic [13] aims to find the subset of variables to destroy \mathcal{X}^t such that it leads to the optimal \mathbf{x}^{t+1} that differs from \mathbf{x}^t on at most k_t variables, i.e., it computes the optimal solution \mathbf{x}^{t+1} that sits within a given Hamming ball of radius k_t centered around \mathbf{x}^t . To find \mathbf{x}^{t+1} , the LB heuristic solves the LB ILP that is exactly the same ILP from input but with one additional constraint that limits the distance between \mathbf{x}^t and \mathbf{x}^{t+1} :

$$\sum_{i \in [n]: x_i^t = 0} x_i^{t+1} + \sum_{i \in [n]: x_i^t = 1} (1 - x_i^{t+1}) \leq k_t.$$

The LB ILP is of the same size of the input ILP (i.e., it has the same number of variables and one more constraint), therefore, it is often slow to run in practice.

3 Related Work

In this section, we summarize related work on LNS for ILPs, LNS-based primal heuristics in BnB and LNS for other COPs.

3.1 LNS for ILPs

While a lot of effort has been made to improve BnB for ILPs in the past decades, LNS for ILPs has not been studied extensively in the past. Recently, Song et al. [40] show that even a randomized destroy heuristic in LNS can outperform state-of-the-art BnB in runtime. In the same paper, they show that an ML-guided decomposition-based LNS can achieve even better performance, where they apply reinforcement learning and imitation learning to learn the destroy heuristics. Since then, there have been a few more recent studies on ML-based LNS for ILPs. Sonnerat et al. [41] learn to select variables to destroy via imitating LB. Wu et al. [43] learn the same thing but they use reinforcement learning instead. The main difference between LB-RELAX and ML-based heuristics is that LB-RELAX does not require extra computational resource for learning and is agnostic to the underlying problem distributions. LB-RELAX also has a better balance between efficiency and effectiveness than those existing non-ML heuristics.

3.2 LNS-Based Primal Heuristics in BnB

LNS-based primal heuristics is one of the rich set of primal heuristics in BnB for ILPs and many techniques have been proposed in past decades. With the same purpose of improving primal bounds of the ILPs, the main differences between the LNS-based primal heuristics in BnB and LNS for ILPs are the following: (1) Since LNS-based primal heuristics are often more expensive to run than the others in BnB, they are executed periodically at different search tree nodes during the main search and the execution schedule is itself dynamic; (2) the destroy heuristics for LNS in BnB are often designed to use information, such

as the dual bound and the LP relaxation at a search tree node, that is specific to BnB and not directly applicable in LNS for ILPs in our setting.

Next, we briefly summarize the destroy heuristics in LNS-based primal heuristics. The Crossover heuristics [37] destroy variables that have different values in a set of selected known solutions (typically two). The Mutation heuristics [37] destroys a random subset of variables. Relaxation Induced Neighborhood Search (RINS) [10] destroys variables whose values disagree in the solution of the LP relaxation at the current search tree node and the current incumbent solution. Relaxation Enforced Neighborhood Search (RENS) [7] restricts the neighborhood to be the feasible roundings of the LP relaxation at the current search tree node. Local Branching [13] restricts the neighborhood to a ball around the current incumbent solution. Distance Induced Neighborhood Search (DINS) [15] takes the intersection of the neighborhoods of the Crossover, LB and RINS heuristics. Graph-Induced Neighborhood Search (GINS) [33] destroys the breadth-first-search neighborhood of a variable in the bipartite graph representation of the ILP. An adaptive LNS primal heuristic that essentially solves a multi armed bandit problem has been proposed to combine the power of these heuristics [18].

LB-RELAX is closely related to RINS [10] since they both use LP relaxations to select neighborhoods. However, RINS is more suitable in BnB since it can adapt dynamically to the constraints added by branching. It uses the LP relaxation of the original problem, whereas LB-RELAX uses that of the LB ILP which takes into account the incumbent solutions that could change from iteration to iteration in LNS.

3.3 LNS for Other COPs

LNS has been applied to solve a wide range of COPs, such as the vehicle routing problem [5, 36], the traveling salesman problem [39], scheduling problems [26, 44] and path planning problems [23, 28, 29]. Recently, ML-based methods have been applied to improve LNS for those applications [9, 20, 22, 30, 32].

4 The Local Branching Relaxation Heuristic

Recently, designing effective destroy heuristics in LNS for ILPs has been a focus in the ML community [40, 41, 43]. However, it is difficult to apply ML-based destroy heuristics to general ILPs since they are often customized for ILPs from certain problem distributions, e.g., graph optimization problems from a given graph distribution or scheduling problems where resources and demands follow the distribution of historical data, and require extra computational resources for training. There has been a lack of study on destroy heuristics that are agnostic to the underlying distribution of the problem. Existing ones such as randomized heuristics are simple and fast but sometimes not effective [40, 41]. LB are effective but not efficient [31, 41] since it exhaustively solves an ILP the same size as input for the best improvement.

Algorithm 2. LB-RELAX (LB-RELAX-S)

```

1: Input: An ILP, incumbent solution  $\mathbf{x}^t$  and neighborhood size  $k_t$ .
2: Construct the LB ILP given  $\mathbf{x}^t$  and  $k_t$ 
3:  $\bar{\mathbf{x}}^{t+1} \leftarrow$  Solve the LP relaxation of the LB ILP
4:  $\Delta_i \leftarrow |\bar{x}_i^{t+1} - x_i^t|$  for all  $i \in [n]$ 
5:  $\bar{\mathcal{X}}^t \leftarrow \{x_i : \Delta_i > 0, i \in [n]\}$ 
6: if  $|\bar{\mathcal{X}}^t| \geq k_t$  then
7:    $\mathcal{X}^t \leftarrow$  Select  $k_t$  variables greedily with the largest  $\Delta_i$  from  $\bar{\mathcal{X}}^t$ 
   ( $\mathcal{X}^t \leftarrow$  Select  $k_t$  variables uniformly at random from  $\bar{\mathcal{X}}^t$ )
8: else
9:    $\mathcal{X}' \leftarrow$  a random subset of  $k_t - |\bar{\mathcal{X}}^t|$  variables from  $\{x_i : \Delta_i = 0, i \in [n]\}$ 
10:   $\mathcal{X}^t \leftarrow \bar{\mathcal{X}}^t \cup \mathcal{X}'$ 
11: return  $\mathcal{X}^t$ 

```

There are well-known approximation algorithms for NP-hard COPs based on LP relaxation [25]. Typically, they solve the LP relaxation of the ILP of the original problem and apply deterministic or randomized rounding afterwards to construct an integral solution. These algorithms often have theoretical guarantee on the effectiveness and are fast, since LP can be solved in polynomial time. Inspired by those algorithms, we propose destroy heuristic LB-RELAX that first solves the LP relaxation of the LB ILP and then constructs the neighborhood (selects variables \mathcal{X}^t to destroy) based on the LP relaxation solution. Specifically, given an ILP and the incumbent solution \mathbf{x}^t in iteration t , we construct the LB ILP with neighborhood size k_t and solve its LP relaxation. Let $\bar{\mathbf{x}}^{t+1}$ be the LP relaxation solution to the LB ILP. Also, let $\Delta_i = |\bar{x}_i^{t+1} - x_i^t|$ and $\bar{\mathcal{X}}^t = \{x_i : \Delta_i > 0, i \in [n]\}$. $\bar{\mathcal{X}}^t$ includes all the fractional variables in the LP relaxation solution and all integral variables that have different values from \mathbf{x}^t . In the following, we introduce (1) LB-RELAX, (2) LB-RELAX-S, a variant of LB-RELAX with randomized sampling and (3) LB-RELAX-R, another variant of LB-RELAX that combines a randomized destroy with LB-RELAX to help avoid local minima more effectively.

LB-RELAX first gets the LP relaxation solution $\bar{\mathbf{x}}^{t+1}$ of the LB ILP and then calculates Δ_i and $\bar{\mathcal{X}}^t$ from $\bar{\mathbf{x}}^{t+1}$, \mathbf{x}^t . To construct \mathcal{X}^t (the set of variables to destroy), it then greedily selects k_t variables with the largest Δ_i and breaks ties uniformly at random. Intuitively, LB-RELAX greedily selects the variables whose values are more likely to change in the incumbent solution \mathbf{x}^t after solving the LB ILP. LB-RELAX is summarized in Algorithm 2. Instead of using the LP relaxation of the LB ILP, one could argue that we alternatively use that of the original ILP similar to RINS [10]. However, the advantage of LB-RELAX over using the LP relaxation of the original problem is that, by approximating the solution to the LB ILP, LB-RELAX selects neighborhoods based on the incumbent solutions that change from iteration to iteration, whereas the original LP relaxation is a static and less informative feature that is pre-computed before the LNS procedure.

LB-RELAX-S is a variant of LB-RELAX with randomized sampling. To construct \mathcal{X}^t , instead of greedily choosing variables with the largest Δ_i , it selects k_t variables from $\bar{\mathcal{X}}^t$ uniformly at random. If $|\bar{\mathcal{X}}^t| < k_t$, it selects all variables from $\bar{\mathcal{X}}^t$ and $k_t - |\bar{\mathcal{X}}^t|$ variables from the remaining uniformly at random. LB-RELAX is summarized in Algorithm 2 where the parts in blue highlight the differences between LB-RELAX and LB-RELAX-S. Since $0 \leq \Delta_i \leq 1$, one could treat Δ_i as a probability distribution and sample k_t variables accordingly (see [41] for an example of how to normalize the distribution to sample k_t variables). However, this variant performs similarly to or slightly worse than LB-RELAX-S empirically and require extra hyperparameter tunings for the normalization. We therefore omit it and focus on the simpler variant in this paper.

LB-RELAX-R is another variant of LB-RELAX that leverages a randomized destroy to avoid local minima more effectively. Once LB-RELAX fails to find an improving solution in iteration t , if we let $k_{t+1} = k_t$, it will solve the exact same LP relaxation of the LB ILP again in the next iteration since the incumbent solution $\mathbf{x}^{t+1} = \mathbf{x}^t$ and the neighborhood size stay the same. Also, since LB-RELAX uses a greedy rule, it will select the same set of variables with the largest Δ_i 's deterministically, except that it might need to break ties randomly in some cases when there are multiple variables with the same Δ_i . Therefore, it is susceptible to getting stuck at local minima. To tackle this issue, once LB-RELAX fails to find a new incumbent solution, we update k_{t+1} using the adaptive method described in the next paragraph. If it fails again in the next iteration, we switch to a randomized destroy heuristic that uniformly samples variables at random without replacement to construct the neighborhood. We switch back to LB-RELAX after running the randomized destroy heuristic for at least γ seconds and a new incumbent solution is found.

Next, we discuss an adaptive method to set the neighborhood size k_t for LB-RELAX and its variants. The initial neighborhood size k_0 is set to a constant or a fraction of the number of variables in the input ILP. In iteration t , if LNS finds a new incumbent solution, we let $k_{t+1} = k_t$. Otherwise, we increase k_t by a factor $\alpha > 1$. Also, we upper bound the neighborhood size k_t to a fraction $\beta < 1$ of the number of variables to make sure the sub-ILP in each iteration is not too difficult to solve, i.e., we let $k_{t+1} = \min\{\alpha \cdot k_t, \beta \cdot n\}$. This adaptive way of choosing k_t also helps address the issue of local minima by expanding the search neighborhood when LNS fails to improve the solution. It is applicable to not only LB-RELAX and its variants but also any destroy heuristics that require a given neighborhood size k_t .

5 Empirical Evaluation

In this section, we demonstrate the efficiency and effectiveness of LB-RELAX and its variants through extensive experiments on ILP benchmarks.

5.1 Setup

Instance Generation. We evaluate on four NP-hard problem benchmarks selected from previous work [38, 40, 43], which consist of synthetic minimum vertex cover (MVC), maximum independent set (MIS), set covering (SC) and multiple knapsack (MK) instances. MVC and MIS instances are generated according to the Barabasi-Albert random graph model [3], with 9,000 nodes and average degree 5 following [40]. SC instances are generated with 4,000 variables and 5,000 constraints following [43]. MK instances are generated with 400 items and 40 knapsacks following [38]. For each problem, we generate 100 instances.

Baselines. We compare LB-RELAX, LB-RELAX-R and LB-RELAX-S with the following baselines:

- BnB using SCIP (v8.0.1) as the solver with the aggressive mode turned on to focus on improving the primal bound;
- LB: LNS which selects the neighborhood with the LB heuristics;
- RANDOM: LNS which selects the neighborhood by uniformly sampling a subset of variables of a given neighborhood size k_t ;
- GRAPH: LNS which selects the neighborhood based on the bipartite graph representation of the ILP similar to GINS [33]. A bipartite graph representation consists of nodes representing the variables and constraints on two sides, respectively, with an edge connecting a variable and a constraint if a variable has a non-zero coefficient in the constraint. It runs a breadth-first search starting from a random variable node in the bipartite graph and selects the first k_t variable nodes expanded.

Furthermore, we compare our approaches with state-of-the-art ML approaches:

- IL-LNS: LNS which selects the neighborhood using a GCN-based policy obtained by learning to imitate the LB heuristic [41]. We implement IL-LNS since the authors do not fully open source the code;
- RL-LNS: LNS which selects the neighborhood using a GCN-based policy obtained by reinforcement learning [43]. Note that this approach does not require a given neighborhood size k_t since the size is defined implicitly by how the trained policy is used. We use the code made available by the authors.

Hyperparameters. We conduct our experiments on 2.5 GHz Intel Xeon Platinum 8259CL CPUs with 32 GB RAM. All experiments use the hyperparameters described below unless stated otherwise. We use SCIP (v8.0.1) [8], the state-of-the-art open source ILP solver for the repair operations in LNS. To run LNS, we find an initial solution by running SCIP for 10s for MVC, MIS and SC and 20s for MK. We set the time limit to 60 min to solve each instance and 2 min for each repair operation in LNS. Except for LB, we set the time limit to 10 min for each repair operation since LB solves a larger ILP than other approaches in each iteration and typically requires a longer time limit. All approaches require a




Fig. 1. Comparison with non-ML approaches: The primal gap as a function of time, averaged over 100 instances.

neighborhood size k_t in LNS, except for BnB and RL-LNS. The initial neighborhood size (k_0) is set to $k_0 = 400, 200, 150$ and 400 for MVC, MIS, SC and MK, respectively. For fair comparison, all baselines use adaptive neighborhood sizes with $\alpha = 1.02$ and $\beta = 0.5$, except for BnB and RL-LNS. For LB-RELAX-R, we set $\gamma = 30$ s. Additional details on tuning hyperparameters are included in Appendix¹.

Metrics. We use the following metrics to evaluate the efficiency and effectiveness of different approaches: (1) The *primal bound* is the objective value of the ILP. (2) The *primal gap* [6] is the normalized difference between the primal bound v and a precomputed best known objective value v^* , defined as $\frac{|v-v^*|}{\max(v, v^*, \epsilon)}$ if v exists and $v \cdot v^* \geq 0$, or 1 otherwise. We use $\epsilon = 10^{-8}$ to avoid division by zero and v^* is the best primal bound found within 60 min by any approach in the portfolio for comparison. (3) The *primal integral* [1] at time q is the integral on $[0, q]$ of the primal gap as a function of time. It captures the quality of and the speed at which solutions are found. (4) The *survival rate* to meet a certain primal gap threshold is the fraction of instances with the primal gap below the threshold [41]. Since BnB and LNS are both anytime algorithms, we

¹ Appendix is available in the full version of the paper: <https://arxiv.org/abs/2212.08183>.




Fig. 2. Comparison with non-ML approaches: The survival rate over 100 instances as a function of time to meet a certain primal gap threshold. The primal gap threshold is chosen from Table 1 as the median of the average primal gaps at 60 min time cutoff over all approaches rounded to the nearest 0.05%.

show the metrics as a function of time or the number of iterations in LNS (when applicable) to demonstrate their anytime performance.

5.2 Results

Comparison with Non-ML Approaches. First, we compare LB-RELAX, LB-RELAX-R and LB-RELAX-S with non-ML approaches, namely BnB, LB, RANDOM and GRAPH. Figure 1 shows the primal gap as a function of time, averaged over 100 instances. The results show that LB-RELAX, LB-RELAX-R and LB-RELAX-S consistently improve the primal gap a lot faster than the baselines in the first few minutes of LNS. LB-RELAX improves the primal gap slightly faster than LB-RELAX-S in all cases. On average, LB-RELAX is always better than the baselines at any point of time on MK instances and LB-RELAX-S is always better than the baselines on SC and MK instances. However, both LB-RELAX and LB-RELAX-S could get stuck at some local minima. In those cases, they need some time to escape local minima by adjusting the neighborhood size and sometimes could be outperformed by some baselines with longer time on the MVC and MIS instances. By adding randomization to LB-RELAX, LB-RELAX-R escapes local minima more efficiently than LB-RELAX and LB-RELAX-S. On

Table 1. Primal gap (PG) (in percent) and primal integral (PI) at 60 min time cutoff, averaged over 100 instances, and their standard deviations.

| | MVC | | MIS | |
|------------|--------------------|--------------------|--------------------|------------------|
| | PG (%) | PI | PG (%) | PI |
| BnB | 1.01 ± 0.46 | 128.6 ± 14.6 | 2.80 ± 1.36 | 144.0 ± 20.1 |
| LB | 0.15 ± 0.08 | 22.1 ± 3.6 | 1.20 ± 0.31 | 56.3 ± 9.4 |
| RANDOM | 0.11 ± 0.05 | 32.3 ± 2.3 | 0.10 ± 0.05 | 18.0 ± 2.5 |
| GRAPH | 0.17 ± 0.04 | 40.8 ± 2.5 | 1.56 ± 0.18 | 90.2 ± 7.6 |
| LB-RELAX | 0.04 ± 0.03 | 10.3 ± 1.7 | 0.39 ± 0.12 | 29.4 ± 4.3 |
| LB-RELAX-R | 0.09 ± 0.04 | 9.6 ± 1.7 | 0.04 ± 0.04 | 9.3 ± 1.7 |
| LB-RELAX-S | 0.42 ± 0.20 | 28.8 ± 8.1 | 0.37 ± 0.11 | 51.7 ± 10.1 |
| | SC | | MK | |
| BnB | 1.15 ± 0.98 | 87.4 ± 38.6 | 0.91 ± 0.59 | 60.7 ± 17.9 |
| LB | 1.23 ± 0.98 | 114.1 ± 35.7 | 1.50 ± 0.48 | 97.7 ± 13.0 |
| RANDOM | 2.68 ± 1.31 | 124.4 ± 45.7 | 1.24 ± 0.36 | 68.9 ± 14.7 |
| GRAPH | 8.75 ± 2.15 | 338.2 ± 77.0 | 0.33 ± 0.14 | 23.6 ± 4.9 |
| LB-RELAX | 1.37 ± 0.96 | 63.9 ± 34.0 | 0.20 ± 0.09 | 11.3 ± 3.0 |
| LB-RELAX-R | 1.14 ± 0.90 | 58.9 ± 31.5 | 0.00 ± 0.00 | 3.7 ± 0.4 |
| LB-RELAX-S | 0.88 ± 0.85 | 63.8 ± 32.4 | 0.19 ± 0.07 | 11.8 ± 2.4 |

average, LB-RELAX-R is always better than the baselines at any point of time in the search on the MVC, MIS and MK instances.

Table 2. The time (in seconds) to improve the initial solution in one iteration and the improvement of the primal bound, averaged over 100 instances. The time for LB is the solving time of the LB ILP. The time for LB-RELAX and LB-RELAX-S is the sum of the solving times of the LB relaxation and the sub-ILP. The numbers in parentheses are the speed-ups. The improvement is computed by taking the difference between the initial solution and the new incumbent solution and the numbers in parentheses are the losses in quality in percent compared to LB. ↑ means higher is better, ↓ means lower is better.

| | | MVC | MIS | SC | MK |
|------------|-------|----------------|---------------|----------------|-----------------|
| LB | Time↓ | 40.2 | 56.0 | 600.0 | 600.0 |
| | Imp.↑ | 129.79 | 65.50 | 12.21 | 216.51 |
| LB-RELAX | Time↓ | 12.1 (3.3x) | 19.5 (2.9x) | 125.3 (4.8x) | 5.87 (102.2x) |
| | Imp.↑ | 129.41 (-0.3%) | 65.19 (-0.5%) | 15.77 (+29.2%) | 141.10 (-34.8%) |
| LB-RELAX-S | Time↓ | 12.0 (3.4x) | 19.5 (2.9x) | 24.51 (24.5x) | 5.12 (117.6x) |
| | Imp.↑ | 128.61 (-0.9%) | 62.46 (-4.6%) | 5.65 (-53.7%) | 113.48 (-47.6%) |

Table 1 presents the average primal gap and primal integral at 60 min time cutoff. (See results at 15, 30 and 45 min time cutoff in Appendix.) On MVC, SC and MK instances, all LB-RELAX, LB-RELAX-S and LB-RELAX-R have lower




Fig. 3. Comparison with LB: The primal bound as a function of the number of iterations, averaged over 100 instances.

primal gaps and primal integrals on average than any baselines, demonstrating that they not only find higher quality solutions but also find them at a faster speed. On MIS and MK instances, LB-RELAX-R achieves the lowest primal gap and primal integral among all approaches. It also achieves the lowest primal integral on MVC and SC instances. Overall, LB-RELAX-R always comes up in the top 2 in both metrics on all problems.

Figure 2 shows the survival rate over 100 instances as a function of time to meet a certain primal gap threshold. On MVC instances, LB-RELAX and LB-RELAX-R achieve final survival rates above 0.9 while the best baseline RANDOM stays below 0.8. On MIS instances, both LB-RELAX-R and RANDOM achieve final survival rates of 1.0 but LB-RELAX-R uses shorter time. On SC instances, LB-RELAX-S and LB-RELAX-R consistently has a higher survival rate than the baselines. On MK instances, LB-RELAX and its variants achieve survival rates above 0.9 within 15 min while the best baseline GRAPH only gets to around 0.6 with 60 min.

One limitation of LB-RELAX and its variants is that they do not perform well on some problem domains, for example the maximum cut and combinatorial auction problems. Please see Appendix for more results.

Next, we run LB, LB-RELAX and LB-RELAX-S for 10 iterations to compare their effectiveness. We follow the same setup as earlier described, except that we do not use adaptive neighborhood sizes to make sure they have the same

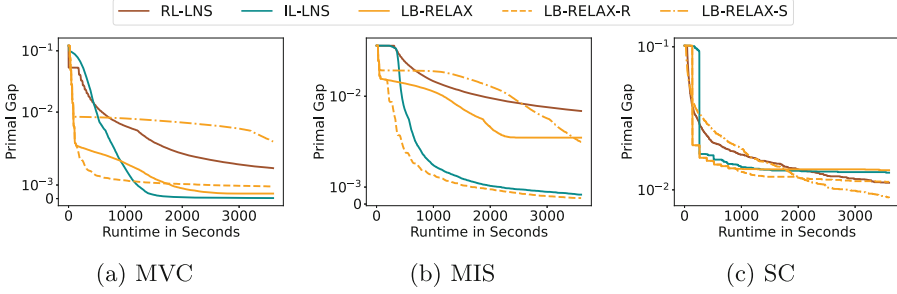


Fig. 4. Comparison with ML approaches: The primal bound as a function of time, averaged over 100 instances.




Fig. 5. Comparison with ML approaches: The survival rate over 100 instances as a function of time to meet a certain primal gap threshold. The primal gap thresholds are chosen in the same way as Fig. 2.

k_t in each iteration t . Note that the time limit for solving the sub-ILP in each iteration is set to 10 min for LB and 2 min for LB-RELAX and LB-RELAX-S. Table 2 shows the average time to improve the initial solutions and the average improvement of the primal bound in the first iteration of LNS. This allows us to compare how closely LB-RELAX and LB-RELAX-S approximate the quality of the neighborhood selected by LB and study the trade-off between quality and time. Compared to LB, LB-RELAX and LB-RELAX-S have 2.9x–117.6x speed-up but only lose at most 53.7% in quality. In particular, on MVC and MIS instances, both LB-RELAX and LB-RELAX-S lose 0.5% to 4.6% in quality but have at least 2.9x speed-up; on SC instances, LB-RELAX even gains 29.2% in quality and save 79.1% in time, due to LB cannot find a good enough neighborhood within its time limit (Fig. 5).

In Fig. 3, we show the primal bound as a function of the number of iterations. It allows comparing the effectiveness of different heuristics independently of their speed. On the MVC instances, both LB-RELAX and LB-RELAX-S perform similarly to but slightly worse than LB. On the SC and MK instances, LB-RELAX achieves better performance than LB, again due to scalability issues of LB, and LB-RELAX-S achieves competitive performance with LB after 10

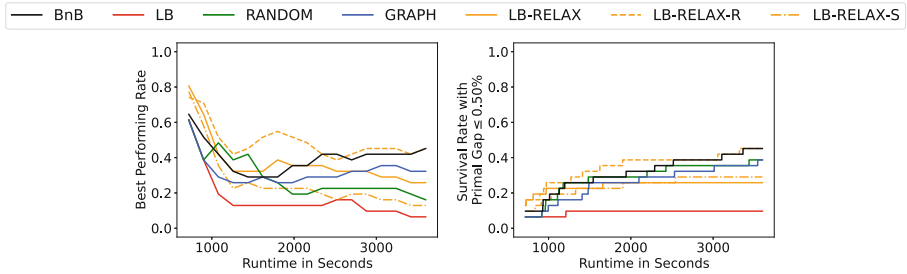


Fig. 6. Results on 31 selected MIPLIB instances: The best performing rate as a function of time (left) and the survival rate over 31 instances as a function of time to meet the primal gap threshold 0.50% (right).

iterations. However on the MIS instances, both LB-RELAX and LB-RELAX-S are able to quickly improve the primal bound in the first 2–3 iterations, but afterwards converge to local minima and the gaps between them and LB increase. To complete the first 10 iterations, both LB-RELAX and LB-RELAX-S take less than 21 min on SC instances and 3.3 min on the others, while LB takes at least 57 min and sometimes up to 100 min.

Comparison with ML Approaches. Then, we compare LB-RELAX, LB-RELAX-R and LB-RELAX-S on MVC, MIS and SC instances with ML approaches, namely IL-LNS and RL-LNS. Figure 4 shows the primal gap as a function of time averaged over 100 instances. The results show that LB-RELAX, LB-RELAX-R and LB-RELAX-S consistently improve the primal bound a lot faster than IL-LNS and RL-LNS in the first few minutes of LNS. On MVC instances, IL-LNS surpasses LB-RELAX-R with the smallest average primal gap best after 20 min and achieve (close-to-)zero gaps after 30 min. On MIS instances, LB-RELAX-R has a smaller gap than both IL-LNS and RL-LNS throughout the first 60 min. On SC instances, IL-LNS is very competitive with LB-RELAX and converges to a similar but slightly higher gap than LB-RELAX-R and LB-RELAX-S; RL-LNS converges to almost the same primal gap as LB-RELAX-R on average but is worse than the best performer LB-RELAX-S. Overall, LB-RELAX and its variants, that do not require extra computational resources for training, are competitive with and more often even better than state-of-the-art ML approaches, suggesting that they are agnostic to the distributions of the instances and easily applicable to different problem domains.

Results on Selected MIPLIB Instances. Finally, we examine how well LB-RELAX and its variants perform on ILPs that are diverse in structures and sizes. We test them on the MIPLIB dataset [16]. MIPLIB contains COPs from various real-world domains. We follow a procedure similar to [43] to filter out instances where we first filter to retain ILP instances with only binary variables. Among these, we select instances that are not too easy to solve but relatively easy to find

a feasible solution for. Specifically, we filter out those that BnB can optimally solve within 3 h (too easy) or BnB cannot find any solutions within 10 min (too hard), which gives us 35 instances. For all LNS approaches, we run BnB for 10 min to find the initial solution and set the time limit to 10 min for each repair operation. The initial neighborhood size k_0 is set to 20% of the number of binary variables. We compare LB-RELAX, LB-RELAX-R and LB-RELAX-S with the non-ML baselines. We further filter out 4 instances that no approach can find a better solution than the initial one, which finally gives us 31 instances.

Figure 6 shows the winning rate as a function of time for each approach on the 31 instances. The best performing rate at a time q for an approach is the fraction of instances on which it achieves the best performance (including ties) compared to all approaches in the portfolio. LB-RELAX, LB-RELAX-R and LB-RELAX-S achieve the best performance with less than 1000 s seconds on 25, 23 and 24 instances out of 35, respectively. LB-RELAX-R has the highest best performing rates at different time cutoffs and ties with BnB at 14 instances at the 60-minute mark. Figure 6 also shows the survival rate over the 31 instances as a function of time to meet the primal gap threshold 0.50%. It demonstrates that RANDOM, GRAPH and BnB are competitive with our approaches but overall LB-RELAX-R has the highest survival rate over time. On some instances, LB-RELAX and its variants can significantly outperform the baselines and we show the anytime performance on those in Appendix.

6 Conclusion

In this paper, we focused on designing effective and efficient destroy heuristics to select neighborhoods in LNS for ILPs. LB is an effective destroy heuristic but is slow to run. We therefore proposed LB-RELAX, LB-RELAX-S and LB-RELAX-R to approximate LB’s decisions by solving its LP relaxation that is a lot faster to run. Empirically, we showed that LB-RELAX, LB-RELAX-S and LB-RELAX-R efficiently selected almost as effective neighborhoods as LB and achieved state-of-the-art performance when compared against non-ML and ML approaches. One limitation of our approaches is that they do not work well on some problem domains, however we showed that they still outperformed the baselines on 14 to 25 (depending on the time cutoff) out of 31 difficult MIPLIB instances that are diverse in problem domains, structures and sizes. The other limitation is that they can get stuck at local minima. To address this issue, we proposed techniques to randomize the heuristics and adaptively adjust the neighborhood sizes. For future work, one could improve LB-RELAX and its variants to make them applicable on more problem domains. In addition, instead of using hard-coded rules for scheduling the randomized heuristic in LB-RELAX-R, one could use adaptive LNS to select destroy heuristics to run. It is also future work to develop theoretical claims to help support and explain the effectiveness of LB-RELAX, LB-RELAX-S, LB-RELAX-R and possibly their other variants.

Acknowledgements. This paper reports on research done while Taoan Huang and Aaron Ferber were interns at Meta AI (FAIR). The research at the University of Southern California was supported by the National Science Foundation (NSF) under grant number 2112533.

References

1. Achterberg, T., Berthold, T., Hendel, G.: Rounding and propagation heuristics for mixed integer programming. In: Klatte, D., Lüthi, H.J., Schmedders, K. (eds.) *Operations Research Proceedings 2011*. *Operations Research Proceedings*, pp. 71–76. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29210-1_12
2. Achterberg, T., Wunderling, R.: Mixed integer programming: analyzing 12 years of progress. In: Jünger, M., Reinelt, G. (eds.) *Facets of Combinatorial Optimization*, pp. 449–481. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38189-8_18
3. Albert, R., Barabási, A.L.: Statistical mechanics of complex networks. *Rev. Mod. Phys.* **74**(1), 47 (2002)
4. Amaral, A.R.: An exact approach to the one-dimensional facility layout problem. *Oper. Res.* **56**(4), 1026–1033 (2008)
5. Azi, N., Gendreau, M., Potvin, J.Y.: An adaptive large neighborhood search for a vehicle routing problem with multiple routes. *Comput. Oper. Res.* **41**, 167–173 (2014)
6. Berthold, T.: Primal heuristics for mixed integer programs. Ph.D. thesis, Zuse Institute Berlin (ZIB) (2006)
7. Berthold, T.: *Rens. Math. Program. Comput.* **6**(1), 33–54 (2014)
8. Bestuzheva, K., et al.: The SCIP optimization suite 8.0. Technical report, Optimization Online (2021). http://www.optimization-online.org/DB_HTML/2021/12/8728.html
9. Chen, X., Tian, Y.: Learning to perform local rewriting for combinatorial optimization. *Adv. Neural Inf. Process. Syst.* **32** (2019)
10. Danna, E., Rothberg, E., Pape, C.L.: Exploring relaxation induced neighborhoods to improve MIP solutions. *Math. Program.* **102**(1), 71–90 (2005)
11. De Vries, S., Vohra, R.V.: Combinatorial auctions: a survey. *INFORMS J. Comput.* **15**(3), 284–309 (2003)
12. Dilkina, B., Gomes, C.P.: Solving connected subgraph problems in wildlife conservation. In: Lodi, A., Milano, M., Toth, P. (eds.) *CPAIOR 2010*. LNCS, vol. 6140, pp. 102–116. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13520-0_14
13. Fischetti, M., Lodi, A.: Local branching. *Math. program.* **98**(1), 23–47 (2003)
14. Gasse, M., Chételat, D., Ferroni, N., Charlin, L., Lodi, A.: Exact combinatorial optimization with graph convolutional neural networks. *Adv. Neural Inf. Process. Syst.* **32** (2019)
15. Ghosh, S.: DINS, a MIP improvement heuristic. In: Fischetti, M., Williamson, D.P. (eds.) *IPCO 2007*. LNCS, vol. 4513, pp. 310–323. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-72792-7_24
16. Gleixner, A., et al.: MIPLIB 2017: data-driven compilation of the 6th mixed-integer programming library. *Math. Program. Comput.* **13**(3), 443–490 (2021). <https://doi.org/10.1007/s12532-020-00194-3>
17. Gurobi Optimization, LLC: Gurobi Optimizer Reference Manual (2022). <https://www.gurobi.com>

18. Hendel, G.: Adaptive large neighborhood search for mixed integer programming. *Math. Program. Comput.* **14**(2), 185–221 (2022)
19. Heragu, S.S., Kusiak, A.: Efficient models for the facility layout problem. *Eur. J. Oper. Res.* **53**(1), 1–13 (1991)
20. Hottung, A., Tierney, K.: Neural large neighborhood search for the capacitated vehicle routing problem. In: *ECAI 2020*, pp. 443–450. IOS Press (2020)
21. Huang, T., Dilkina, B.: Enhancing seismic resilience of water pipe networks. In: *Proceedings of the 3rd ACM SIGCAS Conference on Computing and Sustainable Societies*, pp. 44–52 (2020)
22. Huang, T., Li, J., Koenig, S., Dilkina, B.: Anytime multi-agent path finding via machine learning-guided large neighborhood search. In: *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pp. 9368–9376 (2022)
23. Huang, T., et al.: Deadline-aware multi-agent tour planning. In: *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)* (2023)
24. Khalil, E., Le Bodic, P., Song, L., Nemhauser, G., Dilkina, B.: Learning to branch in mixed integer programming. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 30 (2016)
25. Kleinberg, J., Tardos, E.: *Algorithm Design*. Pearson Education India (2006)
26. Kovacs, A.A., Parragh, S.N., Doerner, K.F., Hartl, R.F.: Adaptive large neighborhood search for service technician routing and scheduling problems. *J. Sched.* **15**(5), 579–600 (2012)
27. Land, A.H., Doig, A.G.: An automatic method for solving discrete programming problems. In: Jünger, M., Liebling, T.M., Naddef, D., Nemhauser, G.L., Pulleyblank, W.R., Reinelt, G., Rinaldi, G., Wolsey, L.A. (eds.) *50 Years of Integer Programming 1958-2008*, pp. 105–132. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-540-68279-0_5
28. Li, J., Chen, Z., Harabor, D., Stuckey, P.J., Koenig, S.: Anytime multi-agent path finding via large neighborhood search. In: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 4127–4135 (2021)
29. Li, J., Chen, Z., Harabor, D., Stuckey, P.J., Koenig, S.: MAPF-LNS2: fast repairing for multi-agent path finding via large neighborhood search. In: *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pp. 10256–10265 (2022)
30. Li, S., Yan, Z., Wu, C.: Learning to delegate for large-scale vehicle routing. *Adv. Neural. Inf. Process. Syst.* **34**, 26198–26211 (2021)
31. Liu, D., Fischetti, M., Lodi, A.: Revisiting local branching with a machine learning lens
32. Lu, H., Zhang, X., Yang, S.: A learning-based iterative method for solving vehicle routing problems. In: *International Conference on Learning Representations* (2019)
33. Maher, S.J., et al.: The SCIP optimization suite 4.0 (2017)
34. Manne, A.S.: On the job-shop scheduling problem. *Oper. Res.* **8**(2), 219–223 (1960)
35. Pohl, I.: Heuristic search viewed as path finding in a graph. *Artif. Intell.* **1**(3–4), 193–204 (1970)
36. Ropke, S., Pisinger, D.: An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. *Transp. Sci.* **40**(4), 455–472 (2006)
37. Rothberg, E.: An evolutionary algorithm for polishing mixed integer programming solutions. *INFORMS J. Comput.* **19**(4), 534–541 (2007)
38. Scavuzzo, L., et al.: Learning to branch with tree MDPS. *arXiv preprint arXiv:2205.11107* (2022)
39. Smith, S.L., Imeson, F.: GLNS: an effective large neighborhood search heuristic for the generalized traveling salesman problem. *Comput. Oper. Res.* **87**, 1–19 (2017)

40. Song, J., Yue, Y., Dilkina, B., et al.: A general large neighborhood search framework for solving integer linear programs. *Adv. Neural. Inf. Process. Syst.* **33**, 20012–20023 (2020)
41. Sonnerat, N., Wang, P., Ktena, I., Bartunov, S., Nair, V.: Learning a large neighborhood search algorithm for mixed integer programs. arXiv preprint [arXiv:2107.10201](https://arxiv.org/abs/2107.10201) (2021)
42. Toth, P., Vigo, D.: *The Vehicle Routing Problem*. SIAM (2002)
43. Wu, Y., Song, W., Cao, Z., Zhang, J.: Learning large neighborhood search policy for integer programming. *Adv. Neural. Inf. Process. Syst.* **34**, 30075–30087 (2021)
44. Žulj, I., Kramer, S., Schneider, M.: A hybrid of adaptive large neighborhood search and tabu search for the order-batching problem. *Eur. J. Oper. Res.* **264**(2), 653–664 (2018)