

LIMIP: Lifelong Learning to Solve Mixed Integer Programs

Sahil Manchanda and Sayan Ranu

Department of Computer Science and Engineering
Indian Institute of Technology, Delhi
{sahil.manchanda,sayanranu}@cse.iitd.ac.in

Abstract

Mixed Integer programs (MIPs) are typically solved by the Branch-and-Bound algorithm. Recently, *Learning to imitate* fast approximations of the expert *strong branching* heuristic has gained attention due to its success in reducing the running time for solving MIPs. However, existing learning-to-branch methods assume that the entire training data is available in a single session of training. This assumption is often not true, and if the training data is supplied in continual fashion over time, existing techniques suffer from *catastrophic forgetting*. In this work, we study the hitherto unexplored paradigm of *Lifelong Learning to Branch* on Mixed Integer Programs. To mitigate catastrophic forgetting, we propose LIMIP, which is powered by the idea of modeling an MIP instance in the form of a bipartite graph, which we map to an embedding space using a bipartite Graph Attention Network. This rich embedding space avoids catastrophic forgetting through the application of knowledge distillation and elastic weight consolidation, wherein we learn the parameters key towards retaining efficacy and are therefore protected from significant drift. We evaluate LIMIP on a series of NP-hard problems and establish that in comparison to existing baselines, LIMIP is up to 50% better when confronted with lifelong learning.

1 Introduction and Related Work

Combinatorial optimization (CO) is a subclass of optimization problems that deals with optimizing a certain objective function by selecting a subset of elements from a finite set. Although CO problems are generally NP-hard (Taha 2014) from a complexity theory viewpoint, still they are solved routinely in diverse fields such as capacity planning, resource allocation, scheduling, and manufacturing (Taha 2014; Manchanda et al. 2020, 2022; Arora, Galhotra, and Ranu 2017; Ranjan et al. 2022; Jain et al. 2021). It is common to formulate most of these CO problems as Mixed integer programs (MIPs) (Achterberg 2007). However, these MIPs are difficult to solve due to the non-convexity of their feasible region. Instead of solving it directly, their LP-relaxed versions are solved (Achterberg 2007). Modern solvers such as SCIP (Gamrath et al. 2020) employ the Branch-and-Bound (B&B) (Achterberg 2007) algorithm to solve these MIPs. B&B recursively partitions the solution




Figure 1: Illustration of catastrophic forgetting on SoA method (Gasse et al. 2019) on Set Cover dataset. The plot shows the rapid increase in the geometric mean of solving time on instances of past problems when the model is further trained *sequentially* on instances of a new problems.

space into a search tree and then prunes subtrees that provably cannot generate an optimal solution. This is an iterative process, which consists of making sequential decisions such as node selection, branching variable selection, etc., to direct the search procedure. The efficiency of the B&B algorithm mainly depends upon the branching variable selection and node selection (Achterberg 2007). In this work, we focus on the former. In modern MIP solvers (Gamrath et al. 2020), the branching variable decisions are generally based upon hard-coded heuristics designed by experts to direct the search process to solve an MIP (Achterberg 2007). Among various heuristics, *strong branching* is one such heuristic that is highly effective in reducing the size of Branch-and-Bound tree (Achterberg 2007; Gasse et al. 2019). However, its main disadvantage is the extremely high computational cost associated with choosing the best variable to branch. Hence, it is rarely used in practice (Achterberg 2007; Huang et al. 2021).

In order to utilize the advantage of the powerful strong branching heuristic, however at a lower computational cost, recently, multiple algorithms have been developed (Gasse et al. 2019; Khalil et al. 2016; Nair et al. 2020; Gupta et al. 2020). At their core, they use *imitation learning* to learn *fast approximations* of the *strong branching* heuristic on a family of MIP instances. These algorithms estimate scores of branching candidate variables quickly on unseen but sim-

ilar MIP instances. This paradigm is popularly known as *learning to branch*. These methods have obtained significant gains in terms of problem solving time over modern MIP solvers such as SCIP.

Despite significant success, existing techniques on learning to branch are limited by the assumption that the entire training data is available in single session of training.

This assumption is not realistic in the context of MIPs as the semantics of CO problems may keep changing with time. Consequently, the data to train the *learning-to-branch* model is dynamic with updates arriving sequentially over time. The above scenario is commonly observed in industries such as shipping and food delivery where entities, such as service locations and warehouses, get added/removed over time. Further, the semantics of the problem such as the demand, supply, customer distribution, facility constraints etc. also fluctuate. Hence, to tackle such scenarios, the model should be capable of learning in an incremental fashion as more data appears over time as retraining the entire model again can be computationally expensive.

The ability to continually learn over time is referred to as *lifelong/continual learning* (Parisi et al. 2019). This aligns with the ability of humans to continually acquire skills throughout their lifespan. Lifelong learning aspires to gain more knowledge sequentially and improve existing model as more data arrives. However, the quintessential failure model of lifelong learning on neural models is *catastrophic forgetting* i.e when new concepts are learned sequentially, the neural model forgets the concepts it learned previously (Kirkpatrick et al. 2017; Parisi et al. 2019)¹.

In this context, we analyze the performance of the state-of-the-art learning-to-branch technique GCNN (Gasse et al. 2019) in the lifelong learning scenario, i.e., training the model sequentially on different problems. In Fig. 1 we observe that the solving time on instances of the SetCover (SC) problem with edge-probability 0.05, i.e., $SC_{0.05}$ increases significantly when the parameters of the model are updated sequentially using training data of $SC_{0.075}$, $SC_{0.125}$ and $SC_{0.2}$ ². Similar phenomenon is observed on other problems in the sequence such as $SC_{0.075}$, etc. It can be clearly concluded from Fig. 1 that GCNN (Gasse et al. 2019) suffers from catastrophic forgetting in the lifelong learning scenario.

In the context of MIPs, a lifelong learning paradigm, due to its nature, promotes efficient learning since retraining from scratch is costly (Parisi et al. 2019), especially on industrial level MIPs where the number of variables and constraints are in orders of millions (Nair et al. 2020). Further, such a paradigm of *continually gaining competencies* on different problems, offers opportunity to transfer the gained knowledge to unseen as well as previously seen problems. Owing to knowledge sharing across problems, these models cope better with low availability of training data. Learning from low-volume training data is important in the context of MIPs since generating training data (state-action pairs)

¹Detailed related work present in Appendix A.1 at <https://arxiv.org/abs/2208.12226>

²Details of the dataset in Experimental Section

itself is a computationally expensive process (Nair et al. 2020). Finally, using a single model that is updated regularly with various competencies, instead of maintaining multiple problem-specific models is desired as it is memory efficient with lower maintenance overhead.

Motivated by the above listed benefits of lifelong learning, in this work we focus on the novel paradigm of *Lifelong Learning to Branch in Mixed Integer Programs*. Our core contributions are as follows.

- **Problem Formulation:** We present the paradigm of Lifelong Learning to Branch in Mixed Integer Programs. To the best of our knowledge, we are the first to investigate this paradigm.
- **Investigation of catastrophic forgetting:** We conduct an empirical investigation of the SoA method (Gasse et al. 2019) and demonstrate that it suffers from catastrophic forgetting when it learns to branch on different problems in succession.
- **Novel Algorithm:** We propose LiMIP, a Lifelong Learning method to solve Mixed Integer Programs. LiMIP encodes the state-space of a problem through an *edge-weighted, bipartite Graph Attention Network*. To mitigate catastrophic forgetting, LiMIP utilizes a novel combination of *Knowledge Distillation* and *Elastic Weight Consolidation* to shield the key parameters of previously learned problems from significant drift.
- **Experimental Evaluation:** We conduct empirical evaluation on a series of NP-hard problems with *drifting* data distribution and evolving constraints. We establish that LiMIP is effective in learning to solve MIPs in a lifelong fashion and overcomes the problem of catastrophic forgetting. Further, it is also capable of transferring the gained knowledge effectively to NP-hard problems with very limited amount of training data.

2 Preliminaries

Definition 1 (Mixed Integer Program). *A mixed-integer linear program is an optimization problem of the form:*

$$\begin{aligned} & \text{minimize } \mathbf{c}^\top \mathbf{x} \\ & \text{subject to } \mathbf{A}\mathbf{x} \leq \mathbf{b} \\ & \mathbf{l} \leq \mathbf{x} \leq \mathbf{u} \\ & \mathbf{x} \in \mathbb{Z}^p \times \mathbb{R}^{n-p} \end{aligned}$$

where n is the total number of variables, p is the number of integral variables. $\mathbf{x} \in \mathbb{Z}^p \times \mathbb{R}^{n-p}$, $\mathbf{A} \in \mathbb{R}^{m \times n}$ is the constraint coefficient matrix, $\mathbf{b} \in \mathbb{R}^m$, right hand side constraint coefficient vector, $\mathbf{c} \in \mathbb{R}^n$ is the objective coefficient vector. Further, the variables $\mathbf{l}, \mathbf{u} \in \mathbb{R}^n$ represent the lower and upper variable bound vectors.

MIPs are solved widely using Branch-and-Bound (B&B) technique, which relaxes the integrality constraints and obtains a continuous linear program (LP). The LP is solved efficiently using the simplex algorithm (Achterberg 2007). In case the relaxed solution is also integral and respects all the constraints, then it is also a solution to the problem (not necessarily optimal). Otherwise B&B decomposes the LP relaxation into two sub-problems, by splitting the feasible

region based upon a variable that does not respect integrality constraints in the current LP solution x^* . Specifically,

$$x_i \leq \lfloor x_i^* \rfloor \vee x_i \geq \lceil x_i^* \rceil, \quad \exists i \leq p \mid x_i^* \notin \mathbb{Z},$$

The B&B solving process repeatedly performs decomposition generating a search tree. The process stops if both the upper and lower bounds are equal or when the feasible regions do not decompose anymore, which is a certificate of infeasibility or optimality. This B&B procedure involves an extremely important step of selecting the fractional *decision* variable to branch upon from the set of candidate variables \mathcal{C} . The chosen variable is used to partition the search space and has a significant impact on the size of the resulting search tree (Achterberg 2007).

Among several heuristics available to choose the branching variable, *strong branching* is widely known to produce the smallest B&B trees. It calculates the expected bound improvement for each candidate variable before performing branching. Although it produces the smallest B&B trees, strong branching requires computing the solution of two LPs for each candidate variable. The cost of finding the best variable is prohibitively high and hence strong branching is not used in practice.

In the B&B setup, the MILP solver is considered to be the *environment*, and the brancher the *agent* (Khalil et al. 2016; Gasse et al. 2019). At the t^{th} decision step, the solver is in state s_t , which comprises the B&B tree with all past branching decisions (Gasse et al. 2019), the best integer solution found so far, the LP solution of each node, the currently focused leaf node, as well as any other solver statistics (for example, the number of times every primal heuristic has been called). In the context of strong branching, at a given state s_t , let a_t be the variable chosen by strong branching among the set of all candidate variables \mathcal{C} . Based upon the above discussion, we now define the problem of *Learning to Branch*.

Problem 1 (Learning to Branch). *For a B&B tree, at the t^{th} decision step of the solver, let the solver be in state s_t and a decision to choose a variable to branch is to be made from a set of candidate variables. Given a collection of state-action (s_t, a_t) pairs obtained from running strong branching, the goal is to learn a scoring function f parameterized by θ that imitates branching decisions made by the strong branching expert.*

Since in our setup we aim to learn on multiple problems, therefore in this context, we refer to each problem as a task.

Problem 2 (Lifelong Learning to Branch). *Given a sequence of tasks $\mathcal{T} = [\mathcal{T}_1, \dots, \mathcal{T}_T]$ of length T , we aim to update the parameters of the model sequentially over time such that at the i^{th} task, when parameter θ_{i-1} is updated to θ_i by training using the instances of the task \mathcal{T}_i , the model avoids catastrophic forgetting on tasks \mathcal{T}_j for $j < i$. Specifically, the increase in running time on problems $\mathcal{T}_j \quad \forall j < i$ using the updated model θ_i should be reasonably low. Additionally, the performance on newly learned tasks should also not be hindered significantly.*

3 LiMIP: Our Proposed Methodology

In this section we describe our proposed method LiMIP. Fig. 2 presents the overview of LiMIP. We first convert a given MIP instance to a bipartite graph. Next, we describe a method to encode the variables and constraints of MIP using an *edge-weighted, bipartite graph attention network (GAT)*. Finally, we describe the procedure of learning the parameters of the model in a lifelong fashion by avoiding catastrophic forgetting. With the outline being set, we next discuss each of these components in detail.

3.1 MIP Representation: State Encoding

Similar to Gasse et al. (2019), to encode the state s_t of the B&B tree at timestep t , we use a bipartite graph representation $G = (\mathbf{V}, \mathbf{E}, \mathbf{C})$. One side of the graph containing n nodes represent the n variables and the other side consisting of m nodes represent the m constraints. There exists an edge between j^{th} variable node and i^{th} constraint node if j^{th} variable appears in the i^{th} constraint. The weight of an edge $e_{i,j}$ corresponds to the value of the coefficient of the variable v_j in the constraint c_i . We use $\mathbf{V} \in R^{n \times d_1}$ to represent variable features, $\mathbf{E} \in R^{n \times m \times 1}$ for edge features, and $\mathbf{C} \in R^{m \times d_2}$ to represent the constraint features. For each of the node in the graph we use the raw solver specific input features of Gasse et al. (2019), which can be found in the appendix A.2. Fig. 2 shows an example of encoding an MIP to a bipartite graph.

3.2 Policy Parameterization: Edge Weighted Bipartite GAT

Observing the weighted and bipartite nature of the graph, it is natural to parameterize the branching variable policy $f_\theta(a|s_t)$ using an *edge-weighted* bipartite GAT. Specifically, for each node in the graph, the attention layer learns to weigh each of the node’s neighbors differently based upon its importance (Vaswani et al. 2017). Since, our graph is bipartite, we perform two levels of message passing through the GAT. Specifically, first we pass message from the variable side to the constraint side to obtain rich representation of the constraint nodes as follows:

$$\mathbf{c}_i = \alpha_{i,i} \theta^C \mathbf{c}_i + \sum_{j \in \mathcal{N}(i)} \alpha_{i,j} \theta^C \mathbf{v}_j \quad (1)$$

Here, \mathbf{c}_i and \mathbf{v}_j refer to the embeddings of i^{th} constraint and j^{th} variable respectively. $\mathcal{N}(i)$ refers to the neighbors of i^{th} node. θ^C refers to MLP associated to constrained side aggregation. α represents the attention coefficient (defined later). Next, we perform message passing from constraint side to variable side. This allows us to generate richer representations for each of the variables nodes.

$$\mathbf{v}_j = \alpha_{j,j} \theta^V \mathbf{v}_j + \sum_{i \in \mathcal{N}(j)} \alpha_{j,i} \theta^V \mathbf{c}_i \quad (2)$$

θ^V refers to weights associated to the variable side aggregation. The attention coefficient α is computed as below:

$$\alpha_{i,j} = \frac{\exp \left(\rho \left((\mathbf{a}^C)^T \left[\theta^C \mathbf{c}_i \parallel \theta^C \mathbf{v}_j \parallel \theta_e^C \mathbf{e}_{i,j} \right] \right) \right)}{\sum_{k \in \mathcal{N}(i)} \exp \left(\rho \left((\mathbf{a}^C)^T \left[\theta^C \mathbf{c}_i \parallel \theta^C \mathbf{v}_k \parallel \theta_e^C \mathbf{e}_{i,k} \right] \right) \right)}$$

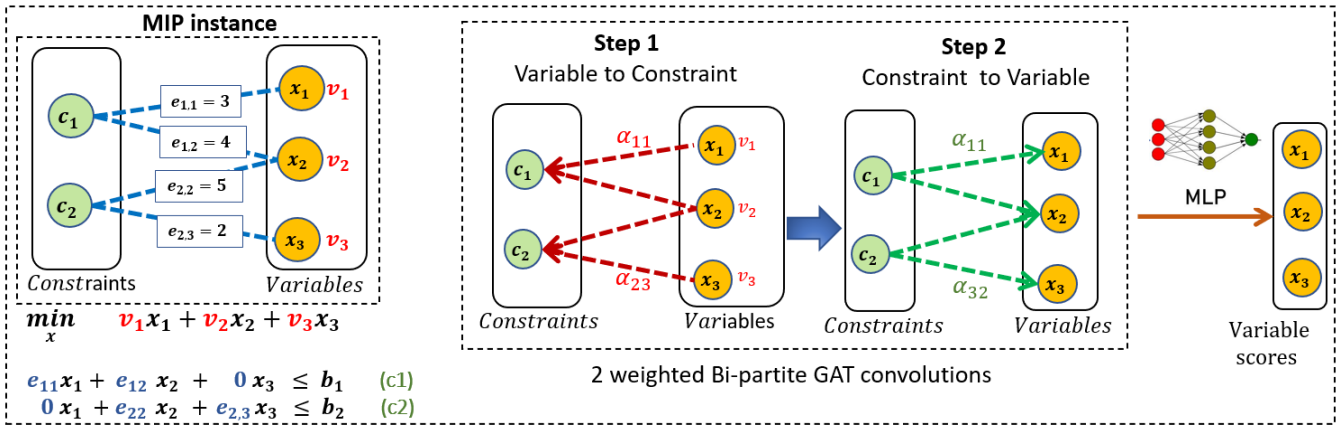


Figure 2: Bipartite graph representation of an MIP with $n = 3$ variables and $m = 2$ constraints. The bipartite graph is encoded via 2 half-aggregations of Bipartite GAT .

The above attention mechanism is parameterized by the weight vector \mathbf{a}^C . θ_e^C on the constraint side aggregation refers to an MLP associated with the edge features. ρ refers to the activation function³. $\alpha_{j,i}$ is defined analogously to $\alpha_{i,j}$ where C is swapped with V and the i^{th} and j^{th} nodes are interchanged.

After the two half-aggregations of eq. 1 and 2, we obtain the final representation of the candidate variable nodes. The final representation of each candidate variable node is passed through a softmax layer to obtain a probability distribution over the variables represented by $f_\theta(a | s_i)$. Further, to stabilize the training procedure of the bipartite-GAT, we use attention mechanism with multiple heads, details of which are present in Appendix Sec.A.3. The detailed architecture is present in Fig. 2.

3.3 Imitation Loss

Since strong branching is a powerful heuristic in reducing the size of the tree, we train the parameters θ by imitation learning of the strong branching rule.

We first collect a set of strong branching state-action pairs $\mathcal{D} = \{(s_k, a_k^*)\}_{k=1}^N$, where N is the number of branching samples collected. Then through imitation learning, we optimize the parameters θ using the following imitation loss function:

$$\mathcal{L}(\theta) = -\frac{1}{N} \sum_{(s, a^*) \in \mathcal{D}} \log f_\theta(a^* | s) \quad (3)$$

The optimization of the above objective encourages the neural model to predict the variable for branching which strong-branching would have chosen.

3.4 Life-Long Learning to Branch

Until now we discussed how to learn the parameters θ of the model for a given task. In this section we discuss how to learn to branch on MIPs in a lifelong fashion. As discussed earlier in Def. 2, we have a set of T problems appearing in sequence $\mathcal{T} = [\mathcal{T}_1, \dots, \mathcal{T}_T]$, and our goal is to learn the parameters sequentially over time where the training data \mathcal{D}_i

³We use LeakyReLU with negative slope = 0.2

for each task \mathcal{T}_i also appears sequentially. A naïve solution is to update the parameters of the model sequentially as new tasks arrive. However, as we already observed in Fig. 1, if the neural model is updated in this fashion, it suffers from catastrophic forgetting on the earlier learned tasks. Hence, our goal is to update the parameters of the model on new tasks while preserving the knowledge gained on previous tasks to avoid catastrophic forgetting.

One way to consolidate past knowledge is to replay the training data of the past tasks. However, as the number of tasks increase, it becomes computationally expensive. Further, another option of storing only a small set of labeled samples and replaying them is prone to over-fitting (Wang et al. 2020). Hence, inspired by recent works on continual learning (Buzzega et al. 2020; Wang et al. 2020), to tackle the problem of catastrophic forgetting in *lifelong learning to branch* we take the following two perspectives. First idea is to approximate the knowledge gained by the model in the past via *distillation* of model’s past behavior when learning new tasks. Second, we optimize the parameters of the model in a constrained way in order to prevent significant drift on the parameters important for previously learned tasks. Fig. 3. visually describes the process. Now, we discuss both the perspectives below in detail.

Mimicking model’s past behavior through Knowledge

Distillation: In order to maintain past learned patterns during lifelong learning, our goal is to search for model parameters that fit well on the current task and also approximate the optimal behavior of the model on the older tasks. Towards this we aim to encourage the model to mimic its original(past) output *logits* for a small number of samples of the past tasks. To accomplish this we apply *Knowledge distillation*(KD) (Buzzega et al. 2020) approach to enforce the neural network to generate similar logits that the model produced for these samples in the past during optimizing of the task to which the related sample belonged to. Mathematically,

$$\mathcal{L}_{KL} = \mathbb{E}_{(s,z) \sim M} [D_{KL}(z || f_\theta(s))] \quad (4)$$

Here $z = f_{\theta_j^*}(s)$ refers to the logits z of sample s and θ_j^* refers to the set of optimal parameters of task \mathcal{T}_j . These

(s, z) pairs are stored in a fixed-size buffer M . Specifically, for $s \in M$, $f_{\theta_j^*}(s)$ is preserved where s is a training sample from task \mathcal{T}_j . When the lifelong learning model is at step i of the sequence \mathcal{T} , M consists of samples of past experiences(logits) for tasks seen till step $i - 1$. Further, since we do not have any prior information of how many tasks we will observe, we use *reservoir sampling* to preserve samples for Knowledge-Distillation. Reservoir sampling ensures that samples from all tasks are stored with equal probability in the buffer without knowing the number of tasks/samples in the stream in advance(Buzzega et al. 2020).

Preservation of model’s important parameters: As we store only a small set of logits in our memory buffer instead of the entire training data of past tasks, it is prone to over-fitting. Although, over-fitting can be tackled to an extent by L2 regularizers, the restriction imposed by L2 regularizers by constraining the entire network through a fixed coefficient is too severe and might prevent learning of the new tasks itself. Inspired from recent works (Kirkpatrick et al. 2017; Wang et al. 2020), to counter this problem, we aim to learn to adjust the magnitude of the parameter updates on certain model weights based on how important they are to the previously learned tasks. To accomplish this we apply *Elastic Weight Consolidation (EWC)* (Wang et al. 2020; Kirkpatrick et al. 2017). Specifically, after the training on a task \mathcal{T}_j is complete, we compute the importance of each parameter w on the task \mathcal{T}_j as follows:

$$\Omega_j^w = \mathbb{E}_{(s, a^*) \sim D_j} \left[\left(\frac{\delta \mathcal{L}(s, a^*)}{\delta \theta_j^w} \right)^2 \right]$$

$\mathcal{L}(s, a^*)$ refers to the loss on sample s with ground-truth a^* . The term $\frac{\delta \mathcal{L}(s, a^*)}{\delta \theta_j^w}$ calculates the gradient of the loss with respect to the parameter w learned on task \mathcal{T}_j . Ω_j^w captures the importance of weight w to task \mathcal{T}_j . We note that D_j is no more required during future tasks once the computation of Ω for task \mathcal{T}_j is complete.

Now, when a new task \mathcal{T}_i arrives, we apply the above regularization (penalize) to prevent large amount of drift on parameters important for earlier learned tasks. Here, the weights of the regularization are obtained from Ω . We accomplish the regularization by the below loss function.

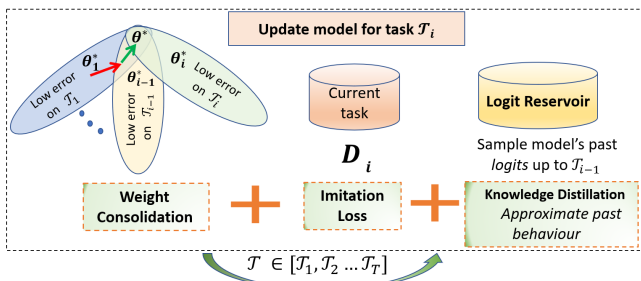


Figure 3: Architecture diagram representing the update mechanism of LiMIP at the i^{th} step of sequence \mathcal{T} .

$$\mathcal{L}_{importance} = \sum_{j=1}^{i-1} \sum_w \Omega_j^w (\theta_i^w - \theta_{j^*}^w)^2 \quad (5)$$

The above term is a quadratic penalty term on the difference between the parameters for the new and the old tasks. Ω consists of diagonal weighing proportional to the diagonal of the Fisher information metric over the old parameters on the old tasks (Liu, Yang, and Wang 2021). θ_{j^*} refers to optimal parameters of task \mathcal{T}_j . When updating parameters of the model to learn to branch on a new task \mathcal{T}_i , the above penalization will encourage the *important model* parameters to be close to the parameters obtained for earlier learned tasks $\mathcal{T}_1, \mathcal{T}_2 \dots \mathcal{T}_{i-1}$. Fig. 2 b) summarizes this concept visually through overlapping optimal parameter spaces.

Lifelong learning optimization objective: Finally, combining the loss functions of eqs. 3, 4 and 5, we obtain the optimization objective at the i^{th} step as follows

$$\begin{aligned} \mathcal{L}_{lifelong} = & \sum_{(s, a^*) \in \mathcal{D}_i} \log f_{\theta} (a^* | s) \\ & + \alpha \mathbb{E}_{(s, z) \sim \mathcal{M}} [D_{KL} (z || f_{\theta}(s))] \\ & + \beta \sum_{j=1}^{i-1} \sum_w \Omega_j^w (\theta_i^w - \theta_{j^*}^w)^2 \end{aligned}$$

The above equation while learning new tasks, consolidates past information in order to maintain stability of parameters important for previously learned tasks. α controls the weight corresponding to mimicking past logits and β controls scale of the weight consolidation regularizer. The detailed steps of training a sequence of tasks in lifelong fashion through LiMIP are described in Algorithm 1 in Appendix A.4.

4 Experiments

In this section we measure the effectiveness of our proposed approach LiMIP and establish:

- **Minimal forgetting:** LiMIP is capable of lifelong learning on NP-hard problems with drifting data distributions and avoids catastrophic forgetting on previously learned problems.
- **No hindrance in learning future tasks:** Despite adding constraints to prevent significant updates to the model, LiMIP does not hinder learning on new tasks.
- **Transfer to Low data regime:** We compare the performance of LiMIP to transfer on a low-training data regime task which is similar to a task LiMIP learned in the past. LiMIP effectively transfers its previously gained and *unforgotten knowledge* to the unseen task.
- **Efficient learning through Bipartite GAT:** Attributed to rich representations learned through the attention mechanism, LiMIP reduces solving time on instances when compared to GCNN (Gasse et al. 2019).

The codebase can be found on <https://github.com/idea-itd/LiMIP>.

4.1 Datasets

We use the following datasets to evaluate the performance of our method against different baselines.

Set Cover: We consider the Set Cover problem of Balas and Ho (1980). Let p be the probability of an item belonging to a set in the Set Cover(SC) problem. SC_p refers to Set Cover problem with set-item probability p . To simulate lifelong learning setup, we generate multiple Set Cover problems datasets each with a different probability, i.e., $\mathcal{T} = [SC_{0.05}, SC_{0.075}, SC_{0.1}, SC_{0.125}, SC_{0.15}, SC_{0.2}]$. In all instances we set number of rows to 700 and number of columns to 800.

Independent Set : We consider the Maximum Independent Set (MIS) problem on the Barabási-Alberta graph (Albert and Barabási 2002) generated with different sizes and affinities. $IS_{A,S}$ denotes as instance where A is the affinity and S is the size of the graph. To simulate lifelong learning setup, we generate independent set problem datasets with different sizes and affinities as $\mathcal{T} = [IS_{4,750}, IS_{4,500}, IS_{4,450}, IS_{5,450}, IS_{5,400}, IS_{5,350}]$.

Facility Location with constraints: We consider the Facility Location problem (Gasse et al. 2019) and to simulate lifelong learning scenario, we use facility capacities and customer demands sampled from drifting distributions over time. The details can be found in Appendix A.5.

These datasets are challenging for state-of-the-art solvers, and also representative of the types of integer programming problems encountered in practice.

4.2 Experimental Setup and Parameters

We use SCIP (Gamrath et al. 2020) as the backend solver, with a time limit of 45 minutes. We use a system running on Intel Xeon 6248 processor with 96 cores and 1 NVIDIA A100 GPU with 40GB memory for our experiments. Similar to existing works (Gasse et al. 2019), we enable cutting planes at the root node and deactivate solver restarts. We keep all other SCIP parameters to default. We use attention mechanism with 2 heads. We set the default buffer size to 500. For details of all parameters and system settings, we refer to App A.6.

Training data generation: For each task, we generate 150,000 branching samples extracted using 10,000 generated instances for training and 30000 validation/test samples generated using 2000 instances.

Metrics: We perform evaluation on 20 different test instances using 5 different SCIP seeds. We report the standard benchmark metric for MILP benchmarking, i.e., the geometric mean of the **running time** of the solver. Additionally, we report the **hardware independent node count** (in Appendix A.8). We compute the average per-instance standard deviation so a value $X \pm s\%$ means it took X secs to solve an instance and while solving one of those instances the time varied by s on an average.

Baselines: We compare our work LiMIP with the state-of-the-art method for learning to branch GCNN (Gasse et al. 2019). We skip comparison with Zarpellon et al. (2020) since it approximates the weaker reliability pseudocost branching heuristic, which has been shown to have

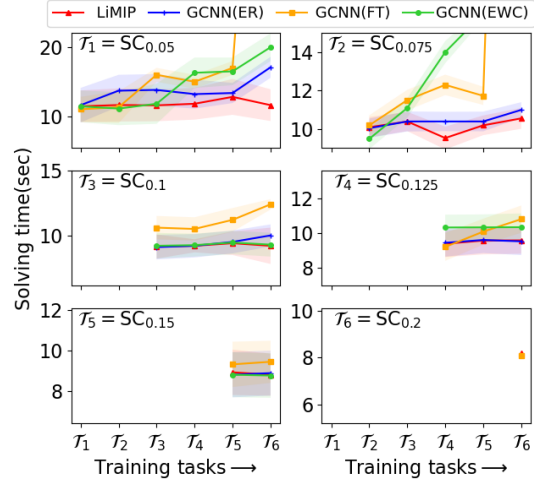


Figure 4: Testing on Set Cover in lifelong scenario: Evolution of solving time for each task when different methods are updated on each task sequentially. Different *evaluation tasks* are shown in *different subplots*. The x-axis denotes the sequence of training tasks and the y-axis denotes the geometric mean of solving time for *test instances* of each task in the sequence. The shaded area refers to standard dev.

an inferior performance in terms of running time (See Appendix A.10). Further, we skip comparison with Gupta et al. (2020) since its focus is on developing CPU based version of learning to branch, which is out of scope of our work. For the sake of completion, we compare with the default SCIP Solver and strong branching in App. A.9.

In the context of lifelong learning, we compare with (1) Fine-tuning (FT) i.e directly updating the model on new tasks as they arrive, (2) Experience Replay (ER) and (3) Elastic Weight Consolidation (EWC) (Kirkpatrick et al. 2017). Details of baseline are present in App. A.7.

4.3 Evaluation in Lifelong Learning Scenario

Evaluating forgetting: In Fig. 4 and 5, in each subplot we study the performance of different methods on test instances of each dataset in the lifelong sequence \mathcal{T} . Specifically, each subplot in these figures refer to a test task and the x-axis shows the sequence of training tasks in the lifelong setup. We observe that as training progresses on different problems in the lifelong setup, the performance of GCNN (FT) (Gasse et al. 2019) on old problem deteriorates significantly. The older the task, the worse is the deterioration. For example, since $SC_{0.05}$ (top, left in Fig. 4) is trained first, it witnesses maximum increase in time across algorithms except the proposed LiMIP. This clearly shows that GCNN suffers from catastrophic forgetting when its parameters are updated on new tasks. In sharp contrast, LiMIP is able to maintain the learned patterns of past tasks when learning new tasks. This is attributed to the *knowledge distillation* loss, which helps in promoting the model to mimic past behavior and the *weight consolidation penalty* term which prevents significant drift on important parameters.

Low hindrance on future tasks: While avoiding catastrophic forgetting is one of the aims of lifelong learning, it should not be at the cost of learning new tasks. From Fig. 4 and Fig. 5, we observe that, while LiMIP does not forget the knowledge it gained in the past, still it does it without impacting future tasks. We can clearly see that on future tasks too LiMIP obtains superior performance compared to existing baselines. In App A.11 we also compare with GAT (FT), GAT (EWC) and GAT (ER). The results on Facility Location Problem are present in App A.9.

4.4 Transferability on Low Data Regimes

As a model learns on more tasks and gains several competencies, it can be utilized as a *weight initializer* to learn on an unseen task with low availability of training data, especially in the case of strong branching where obtaining training data is costly. To test the performance of LiMIP on transfer learning, we create a new dataset $SC_{0.047}$ with only 300 branching samples for training. This number is extremely low compared to other datasets where number of samples = 150000. We perform lifelong learning on $\mathcal{T} = [SC_{0.05}, SC_{0.075}, SC_{0.1}, SC_{0.125}, SC_{0.15}, SC_{0.2}]$ and then fine-tune on $SC_{0.047}$. We compare it with training a model from scratch on the 300 branching samples of $SC_{0.047}$. In Table 1, we study the performance gain obtained using a fine-tuned lifelong learned model vs. a model learned from scratch (GCNN $SC_{0.047}$). We observe that on an average the running time of the LiMIP method fine-tuned on $SC_{0.047}$ model is significantly better than the model trained from scratch.

4.5 Ablation Studies

GCN vs Bipartite GAT: In this section we study the impact of using our Bipartite GAT compared to mean pool based GCNN. In Table 2 we observe that Bipartite GAT improves over GCNN by a small margin in terms of both running time and number of tree nodes.

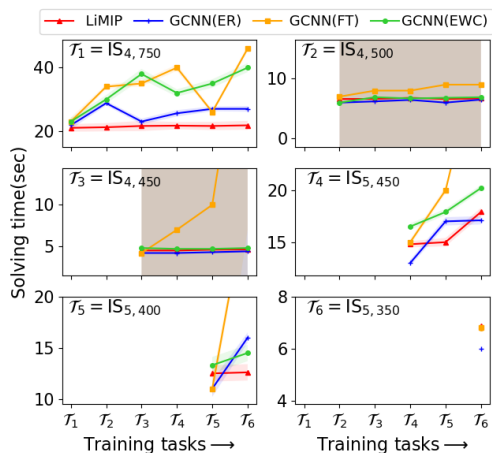


Figure 5: Testing on Independent Set in lifelong scenario: Evolution of solving time for each task when different methods are updated on each task sequentially.

Model	Time(sec)	# Nodes
GCNN $SC_{0.047}$	16.01	607
Bipartite GAT $SC_{0.047}$	17.06	702
Bipartite GAT(FT) + $SC_{0.047}$	16.07	702
LiMIP	15.05	602
LiMIP + ($SC_{0.047}$)	14.05	441

Table 1: Transferability performance: Test performance comparison of fine-tuned model against model trained from scratch on the $SC_{0.047}$ dataset. LiMIP and Bipartite GAT (FT) were trained sequentially on $[SC_{0.05}, SC_{0.075}, SC_{0.1}, SC_{0.125}, SC_{0.15}, SC_{0.2}]$ and then fine-tuned on $SC_{0.047}$

Dataset	Method	Time(sec)	Nodes
$SC_{0.2}$	Bipartite GAT	7.9 ± 2.56	87.2 ± 15.31
	GCNN	8.21 ± 2.45	91.2 ± 13.61
$IS_{4,750}$	Bipartite GAT	22.25 ± 1.24	555.6 ± 6.02
	GCNN	25.73 ± 1.32	672.2 ± 6.2
$FC_{(40,50),(5,10)}$	Bipartite GAT	33.94 ± 1.11	246.10 ± 4.02
	GCNN	35.14 ± 1.30	248.20 ± 4.70

Table 2: Performance comparison between Bipartite GAT and GCNN encoding.

Impact of regularizer and buffer size on lifelong learning: In App A.11 we study the impact of weight regularization and buffer size.

5 Conclusion

Learning-to-Branch techniques have shown significant success in reducing the solving time of Mixed Integer Programs. Although, significant progress has been made, the paradigm of *learning to branch in a lifelong fashion* was unexplored. In this work we first examined the behavior of existing techniques in the lifelong learning scenario and discovered that they suffer from catastrophic forgetting. To tackle this problem, in this work we study the hitherto unexplored paradigm of *Lifelong Learning to Branch* on Mixed Integer Programs. We propose a method LiMIP powered by a Bipartite GAT to encode MIP instances. Further, to mitigate catastrophic forgetting, we apply *knowledge distillation* and *elastic weight consolidation* to shield key parameters from drifting and thereby retaining efficacy. Through extensive experiments on multiple NP-hard problems, we established that LiMIP is able to mitigate forgetting significantly better compared to existing baselines when confronted with lifelong learning. Additionally, the proposed method does not hinder the performance on future learning tasks too.

Acknowledgements

We acknowledge Qualcomm for supporting this research. We acknowledge Google for supporting Sahil Manchanda for this travel.

References

- Achterberg, T. 2007. Constraint Integer Programming. *Ph. D. Thesis, Technische Universität Berlin*.
- Albert, R.; and Barabási, A.-L. 2002. Statistical mechanics of complex networks. *Reviews of modern physics*, 74(1): 47.
- Arora, A.; Galhotra, S.; and Ranu, S. 2017. Debunking the myths of influence maximization: An in-depth benchmarking study. In *Proceedings of the 2017 ACM international conference on management of data*, 651–666.
- Balas, E.; and Ho, A. 1980. Set covering algorithms using cutting planes, heuristics, and subgradient optimization: a computational study. In *Combinatorial Optimization*, 37–60. Springer.
- Buzzega, P.; Boschini, M.; Porrello, A.; Abati, D.; and Calderara, S. 2020. Dark experience for general continual learning: a strong, simple baseline. *Advances in neural information processing systems*, 33: 15920–15930.
- Gamrath, G.; Anderson, D.; Bestuzheva, K.; Chen, W.-K.; Eifler, L.; Gasse, M.; Gemander, P.; Gleixner, A.; Gottwald, L.; and Halbig, K. 2020. The SCIP Optimization Suite 7.0. *ZIB-Report*.
- Gasse, M.; Chételat, D.; Ferroni, N.; Charlin, L.; and Lodi, A. 2019. Exact combinatorial optimization with graph convolutional neural networks. *Advances in Neural Information Processing Systems*, 32.
- Gupta, P.; Gasse, M.; Khalil, E.; Mudigonda, P.; Lodi, A.; and Bengio, Y. 2020. Hybrid models for learning to branch. *Advances in neural information processing systems*, 33: 18087–18097.
- Huang, L.; Chen, X.; Huo, W.; Wang, J.; Zhang, F.; Bai, B.; and Shi, L. 2021. Branch and bound in mixed integer linear programming problems: A survey of techniques and trends. *arXiv preprint arXiv:2111.06257*.
- Jain, J.; Bagadia, V.; Manchanda, S.; and Ranu, S. 2021. NeuroMLR: Robust and Reliable Route Recommendation on Road Networks. *Advances in Neural Information Processing Systems*, 34: 22070–22082.
- Khalil, E.; Le Bodic, P.; Song, L.; Nemhauser, G.; and Dilkina, B. 2016. Learning to branch in mixed integer programming. In *AAAI*, volume 30.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; Hassabis, D.; Clopath, C.; Kumaran, D.; and Hadsell, R. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13): 3521–3526.
- Liu, H.; Yang, Y.; and Wang, X. 2021. Overcoming catastrophic forgetting in graph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 8653–8661.
- Manchanda, S.; Michel, S.; Drakulic, D.; and Andreoli, J.-M. 2022. On the Generalization of Neural Combinatorial Optimization Heuristics. *arXiv preprint arXiv:2206.00787*.
- Manchanda, S.; Mittal, A.; Dhawan, A.; Medya, S.; Ranu, S.; and Singh, A. 2020. GCOMB: Learning Budget-constrained Combinatorial Algorithms over Billion-sized Graphs. *Advances in Neural Information Processing Systems*, 33.
- Nair, V.; Bartunov, S.; Gimeno, F.; von Glehn, I.; Lichocki, P.; Lobov, I.; O’Donoghue, B.; Sonnerat, N.; Tjandraatmadja, C.; Wang, P.; et al. 2020. Solving mixed integer programs using neural networks. *arXiv preprint arXiv:2012.13349*.
- Parisi, G. I.; Kemker, R.; Part, J. L.; Kanan, C.; and Wermter, S. 2019. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113: 54–71.
- Ranjan, R.; Grover, S.; Medya, S.; Chakaravarthy, V.; Sabharwal, Y.; and Ranu, S. 2022. GREED: A Neural Framework for Learning Graph Distance Functions. In Oh, A. H.; Agarwal, A.; Belgrave, D.; and Cho, K., eds., *NeurIPS*.
- Taha, H. A. 2014. *Integer programming: theory, applications, and computations*. Academic Press.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, J.; Song, G.; Wu, Y.; and Wang, L. 2020. Streaming graph neural networks via continual learning. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, 1515–1524.
- Zarpellon, G.; Jo, J.; Lodi, A.; and Bengio, Y. 2020. Parameterizing branch-and-bound search trees to learn branching policies. *arXiv preprint arXiv:2002.05120*.