
Exploring the Power of Graph Neural Networks in Solving Linear Optimization Problems

Chendi Qian
RWTH Aachen University

Didier Chételat
Huawei Noah’s Ark Lab

Christopher Morris
RWTH Aachen University

Abstract

Recently, machine learning, particularly message-passing graph neural networks (MPNNs), has gained traction in enhancing exact optimization algorithms. For example, MPNNs speed up solving mixed-integer optimization problems by imitating computational intensive heuristics like strong branching, which entails solving multiple linear optimization problems (LPs). Despite the empirical success, the reasons behind MPNNs’ effectiveness in emulating linear optimization remain largely unclear. Here, we show that MPNNs can simulate standard interior-point methods for LPs, explaining their practical success. Furthermore, we highlight how MPNNs can serve as a lightweight proxy for solving LPs, adapting to a given problem instance distribution. Empirically, we show that MPNNs solve LP relaxations of standard combinatorial optimization problems close to optimality, often surpassing conventional solvers and competing approaches in solving time.

1 INTRODUCTION

Recently, there has been a surge of interest in training *message-passing graph neural networks* (MPNNs) to imitate steps of classical algorithms, such as for shortest-path problems Cappart et al. (2021); Veličković et al. (2020). As many of those problems can be formulated as linear optimization problems (LPs), it is natural to ask whether MPNNs could be trained to solve general LPs, at least approximately.

Another recent line of research makes this question particularly intriguing. In integer linear optimization,

state-of-the-art solvers all rely on the branch-and-bound algorithm, in which one must repeatedly select variables, subdividing the search space. The best-known heuristic for variable selection is known as “strong branching,” which entails solving LPs to score the variables. This heuristic is, unfortunately, too computationally expensive to use in practice. However, in recent years, there has been a collection of works (Gasse et al., 2019; Gupta et al., 2020, 2022; Nair et al., 2020; Seyfi et al., 2023) that have proposed to use MPNNs to imitate strong branching, with impressive empirical success. No theoretical explanation has ever been put forward to explain this success. However, perhaps the most straightforward explanation would be that MPNNs implicitly learn to imitate the LP solving underlying strong branching.

Chen et al. (2023) provide a first step towards an explanation. In this work, the authors propose to encode LPs as bipartite graphs in the spirit of (Gasse et al., 2019) and show that MPNNs, in principle, can learn to predict the optimal solution of an LP instance up to arbitrary small ε concerning the supremum norm. They also provide some small-scale experiments suggesting that MPNNs can learn to approximate LP solutions surprisingly well. While a step in the right direction, their theoretical result heavily relies on invoking the universal approximation theorem for multi-layer perceptrons Cybenko (1992); Leshno et al. (1993), and therefore does not explain why modestly-sized MPNNs could be particularly effective LP solvers in practice.

In this paper, we present instead a more specific explanation. We show that several variants of *interior-point methods* (IPMs) Gondzio (2012); Nocedal and Wright (2006), an established class of polynomial-time algorithms for solving LPs, can be interpreted as an MPNN with a specific architecture and choice of parameters that takes as input a graph representation of the LP. Specifically, for two common IPM variants, we show that a sequence of standard MPNN steps can emulate a single iteration of the algorithm on a tripartite (rather than bipartite) graph representation of the LP. This novel theoretical result suggests several conclu-

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

sions. First, it indicates that MPNNs appear successful at imitating LP solving because they might often be imitating IPMs and that there is a close connection between MPNNs and this specific LP algorithm. Secondly, although the MPNN’s architecture in our theoretical result involves many MPNN layers, MPNNs are a natural choice of machine learning model for LP solving. It is likely possible to solve LPs approximately with fewer layers.

To prove the last hypothesis, we train MPNNs with several layers less than predicted by our theoretical results to imitate the output of a practical IPM algorithm for LP solving, resulting in the *IPM-MPNN* architecture. Our empirical results show that IPM-MPNNs can lead to reduced solving times compared to a state-of-the-art LP solver with time constraints and competing neural network-based approaches; see Figure 1b for an overview of our approach.

In summary, our findings significantly contribute to the theoretical framework of data-driven exact optimization using MPNNs, and we also showcase the potential of MPNNs in serving as a light-weight proxy for solving LPs.

1.1 Additional related works

In the following, we discuss relevant related work.

MPNNs MPNNs (Gilmer et al., 2017; Scarselli et al., 2009) have emerged as a flexible framework for machine learning on graphs and relational data. Notable instances of this architecture include, e.g., Duvenaud et al. (2015); Hamilton et al. (2017); Veličković et al. (2018), and the spectral approaches proposed in, e.g., Bruna et al. (2014); Defferrard et al. (2016); Kipf and Welling (2017)—all of which descend from early work in Baskin et al. (1997); Kireev (1995); Merkwirth and Lengauer (2005); Micheli (2009); Micheli and Sestito (2005); Scarselli et al. (2009); Sperduti and Starita (1997).

Machine learning for combinatorial optimization Bengio et al. (2021); Kotary et al. (2021) discuss and review machine learning approaches to enhance combinatorial optimization (CO). Concrete examples include the imitation of computationally intensive variable selection rules within the branch-and-cut framework (Khalil et al., 2016; Zarpellon et al., 2020), learning to run (primal) heuristics (Khalil et al., 2017; Chmiela et al., 2021), learning decompositions of large MILPs for scalable heuristic solving (Song et al., 2020), learning to generate cutting planes (Deza and Khalil, 2023) or leveraging machine learning to find (primal) solutions to stochastic integer problems quickly (Bengio et al., 2020).

MPNNs for CO Many prominent CO problems involve graph or relational structures, either directly given as input or induced by the variable-constraint interactions. Recent progress in using MPNNs to bridge the gap between machine learning and combinatorial optimization is surveyed in Cappart et al. (2021).

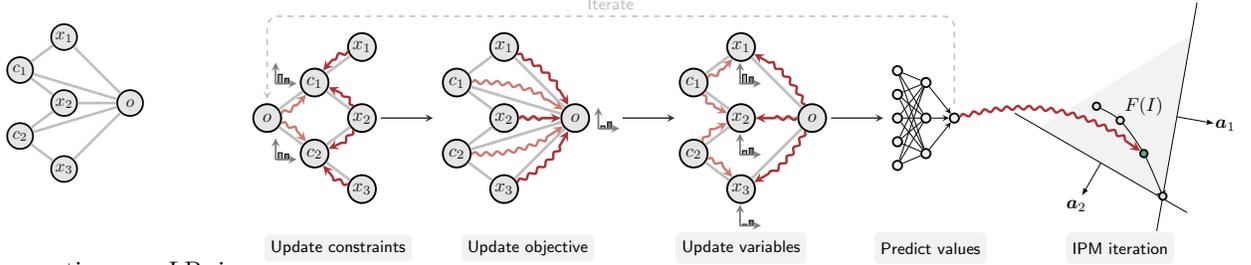
Most relevant to the present work, Gasse et al. (2019) proposed to encode the variable-constraint interaction of a mixed-integer linear optimization problem as a bipartite graph and trained MPNNs in a supervised fashion to imitate the costly strong branching heuristic, which entails solving multiple linear optimization problems, within the branch-and-cut framework (Achterberg et al., 2005). Building on that, Gupta et al. (2020) proposed a hybrid branching model using an MPNN at the initial decision point and a light multi-layer perceptron for subsequent steps, showing improvements on pure CPU machines. Subsequently, Nair et al. (2020) expanded the MPNN approach to branching by implementing a GPU-friendly parallel linear programming solver using the alternating direction method of multipliers that allows scaling the strong branching expert to substantially larger instances, also combining this innovation with a novel MPNN approach to diving. For all the above works, it remains largely unclear why MPNNs are good at (approximately) predicting strong branching scores. Moreover, Khalil et al. (2022) used MPNNs to predict the probability of variables being assigned to 0 or 1 in near-optimal solutions of binary-integer linear optimization problems.

Ding et al. (2020) used MPNNs on a tripartite graph consisting of variables, constraints, and a single objective node enriched with hand-crafted node features. The target is to predict the 0-1 values of the so-called *stable variables*, i.e., variables whose assignment does not change over a set of pre-computed feasible solutions. Li et al. (2022) used MPNNs on a bipartite graph together with a pointer network (Bello et al., 2016) to reorder the variables of a given LP instance, resulting in reduced solving time. Fan et al. (2023) leveraged MPNNs to find good initial basis solutions for the Simplex algorithm.

Finally Wu and Lissner (2023) expressed an LP as an ordinary differential equations system whose state solution converges to the LP’s optimal solution and trained a feed-forward neural network to approximate this state solution. However, their approach hinges on the need to compute the Jacobian matrix, rendering its training phase computationally costly.

2 BACKGROUND

In the following, we describe the necessary background.



(a) Representing an LP instance with two constraints and three variables as a tripartite graph.

(b) IPM-MPNNs emulate interior-point methods.

Figure 1: Overview of our IPM-MPNN framework.

Notation Let $\mathbb{N} := \{1, 2, 3, \dots\}$. For $n \geq 1$, let $[n] := \{1, \dots, n\} \subset \mathbb{N}$. We use $\{\!\{ \dots \}\!\}$ to denote multisets, i.e., the generalization of sets allowing for multiple instances for each of its elements. A *graph* G is a pair $(V(G), E(G))$ with *finite* sets of *vertices* or *nodes* $V(G)$ and *edges* $E(G) \subseteq \{\{u, v\} \subseteq V(G) \mid u \neq v\}$. An *attributed graph* G is a triple $(V(G), E(G), a)$ with a graph $(V(G), E(G))$ and (vertex-)attribute function $a: V(G) \rightarrow \mathbb{R}^{1 \times d}$, for some $d > 0$. Then $a(v)$ are a *node attributes* or *features* of v , for v in $V(G)$. Equivalently, we define an n -vertex attributed graph $G := (V(G), E(G), a)$ as a pair $\mathbf{G} = (G, \mathbf{L})$, where $G = (V(G), E(G))$ and \mathbf{L} in $\mathbb{R}^{n \times d}$ is a *node attribute matrix*. Here, we identify $V(G)$ with $[n]$. For a matrix \mathbf{L} in $\mathbb{R}^{n \times d}$ and v in $[n]$, we denote by \mathbf{L}_v in $\mathbb{R}^{1 \times d}$ the v th row of \mathbf{L} such that $\mathbf{L}_v := a(v)$. We also write \mathbb{R}^d for $\mathbb{R}^{d \times 1}$. The *neighborhood* of v in $V(G)$ is denoted by $N(v) := \{u \in V(G) \mid (v, u) \in E(G)\}$. Moreover, let $\mathbf{x} \in \mathbb{R}^{1 \times d}$, then $\mathbf{D}(\mathbf{x})$ denotes the diagonal matrix with diagonal \mathbf{x} , $\mathbf{0}$ and $\mathbf{1}$ denote the vector of zero and ones, respectively, with an appropriate number of entries. By default, a vector $\mathbf{x} \in \mathbb{R}^d$ is a column vector.

Linear optimization problems A *linear optimization problem* (LP) aims at optimizing a linear function over a feasible set described as the intersection of finitely many half-spaces, i.e., a polyhedron. We restrict our attention to feasible and bounded LPs. Formally, an instance I of an LP is a tuple $(\mathbf{A}, \mathbf{b}, \mathbf{c})$, where \mathbf{A} is a matrix in $\mathbb{Q}^{m \times n}$, and \mathbf{b} and \mathbf{c} are vectors in \mathbb{Q}^m and \mathbb{Q}^n , respectively. We aim at finding a vector \mathbf{x}^* in \mathbb{Q}^n that minimizes $\mathbf{c}^\top \mathbf{x}^*$ over the *feasible set*

$$F(I) = \{\mathbf{x} \in \mathbb{Q}^n \mid \mathbf{A}_j \mathbf{x} \leq b_j \text{ for } j \in [m] \text{ and } x_i \geq 0 \text{ for } i \in [n]\}. \quad (1)$$

In practice, LPs are solved using the Simplex method or polynomial-time IPMs (Nocedal and Wright, 2006).

We now detail the theoretical result that motivates our approach. We first summarize both interior point

methods for linear optimization and MPNNs and then prove a theorem that relates the two.

Interior-point methods for linear optimization

IPMs are algorithms for solving constrained optimization problems. They are particularly efficient for linear optimization, where they were first developed as a (polynomial-time) alternative to the Simplex methods. Variants of the algorithm differ in theoretical guarantees and empirical performance but revolve around the same core approach (Shanno, 2012). In short, the LP to solve is replaced by a perturbed family of problems where a barrier penalty has replaced hard constraints with a parameter $\mu > 0$. IPMs alternate between taking a Newton step to solve this perturbed problem and decreasing this parameter $\mu > 0$, eventually converging to the optimal solution of the original problem.

For concreteness, we present two variants of the approach, an algorithm with theoretical guarantees and a practical algorithm that could be used in practice, following Nocedal and Wright (2006, Chapter 14) and Gondzio (2012), respectively. In the next section, we will show that both algorithms can be connected to MPNNs.

The core idea of IPMs is as follows. First, we consider a perturbed version of the LP (1),

$$\min_{\mathbf{x} \in \mathbb{Q}^n} \mathbf{c}^\top \mathbf{x} - \mu [\mathbf{1}^\top \log(\mathbf{b} - \mathbf{A}\mathbf{x}) + \mathbf{1}^\top \log(\mathbf{x})] \quad (2)$$

for some $\mu > 0$. By introducing the variables $s_i = \mu/x_i$, $r_j = \mathbf{A}_j \mathbf{x} - b_j$ and $w_j = \mu/r_j$, the first-order optimality conditions for (2) can be written as a system

$$\begin{aligned} \mathbf{A}\mathbf{x}^* - \mathbf{r}^* &= \mathbf{b} \\ \mathbf{A}^\top \mathbf{w}^* + \mathbf{s}^* &= \mathbf{c} \\ x_i^* s_i^* &= \mu & i \in [n], \\ w_i^* r_i^* &= \mu & j \in [m], \end{aligned}$$

with $\mathbf{x}^*, \mathbf{w}^*, \mathbf{s}^*, \mathbf{r}^* \geq \mathbf{0}$. Let $\sigma \in (0, 1)$ be another hyperparameter. The two algorithms start from an initial positive point $(\mathbf{x}_0, \mathbf{w}_0, \mathbf{s}_0, \mathbf{r}_0) > \mathbf{0}$, and alternate between computing the Newton step for the perturbed problem (2) at barrier parameter $\sigma\mu$

$$\begin{bmatrix} \mathbf{A} & \mathbf{0} & \mathbf{0} & -\mathbf{I} \\ \mathbf{0} & \mathbf{A}^\top & \mathbf{I} & \mathbf{0} \\ \mathbf{D}(\mathbf{s}) & \mathbf{0} & \mathbf{D}(\mathbf{x}) & \mathbf{0} \\ \mathbf{0} & \mathbf{D}(\mathbf{r}) & \mathbf{0} & \mathbf{D}(\mathbf{w}) \end{bmatrix} \begin{bmatrix} \Delta\mathbf{x} \\ \Delta\mathbf{w} \\ \Delta\mathbf{s} \\ \Delta\mathbf{r} \end{bmatrix} = \begin{bmatrix} \mathbf{b} - \mathbf{A}\mathbf{x} + \mathbf{r} \\ \mathbf{c} - \mathbf{A}^\top\mathbf{w} - \mathbf{s} \\ \sigma\mu\mathbf{1} - \mathbf{D}(\mathbf{x})\mathbf{D}(\mathbf{s})\mathbf{1} \\ \sigma\mu\mathbf{1} - \mathbf{D}(\mathbf{w})\mathbf{D}(\mathbf{r})\mathbf{1} \end{bmatrix},$$

and taking a step in that direction with length $\alpha > 0$, such that the resulting point $(\mathbf{x}', \mathbf{w}', \mathbf{s}', \mathbf{r}') = (\mathbf{x}, \mathbf{w}, \mathbf{s}, \mathbf{r}) + \alpha(\Delta\mathbf{x}, \Delta\mathbf{w}, \Delta\mathbf{s}, \Delta\mathbf{r})$ satisfies $(\mathbf{x}', \mathbf{w}', \mathbf{s}', \mathbf{r}') > \mathbf{0}$.

The above system can be simplified as follows. First, we can infer that

$$\Delta\mathbf{s} = \sigma\mu\mathbf{D}(\mathbf{x})^{-1}\mathbf{1} - \mathbf{s} - \mathbf{D}(\mathbf{x})^{-1}\mathbf{D}(\mathbf{s})\Delta\mathbf{x}, \quad (3)$$

$$\Delta\mathbf{r} = \sigma\mu\mathbf{D}(\mathbf{w})^{-1}\mathbf{1} - \mathbf{r} - \mathbf{D}(\mathbf{w})^{-1}\mathbf{D}(\mathbf{r})\Delta\mathbf{w}, \quad (4)$$

which implies that

$$\begin{aligned} \mathbf{A}\Delta\mathbf{x} + \mathbf{D}(\mathbf{w})^{-1}\mathbf{D}(\mathbf{r})\Delta\mathbf{w} &= \mathbf{b} - \mathbf{A}\mathbf{x} + \sigma\mu\mathbf{D}(\mathbf{w})^{-1}\mathbf{1}, \\ \mathbf{A}^\top\Delta\mathbf{w} - \mathbf{D}(\mathbf{x})^{-1}\mathbf{D}(\mathbf{s})\Delta\mathbf{x} &= \mathbf{c} - \mathbf{A}^\top\mathbf{w} - \sigma\mu\mathbf{D}(\mathbf{x})^{-1}\mathbf{1}. \end{aligned}$$

Therefore, $\Delta\mathbf{x} =$

$$\mathbf{D}(\mathbf{s})^{-1}\mathbf{D}(\mathbf{x})[\mathbf{A}^\top\Delta\mathbf{w} - \mathbf{c} + \mathbf{A}^\top\mathbf{w} + \sigma\mu\mathbf{D}(\mathbf{x})^{-1}\mathbf{1}] \quad (5)$$

$$\begin{aligned} \mathbf{Q}\Delta\mathbf{w} &= \mathbf{b} - \mathbf{A}\mathbf{x} + \sigma\mu\mathbf{D}(\mathbf{w})^{-1}\mathbf{1} \\ &+ \mathbf{A}\mathbf{D}(\mathbf{s})^{-1}\mathbf{D}(\mathbf{x})[\mathbf{c} - \mathbf{A}^\top\mathbf{w} - \sigma\mu\mathbf{D}(\mathbf{x})^{-1}\mathbf{1}] \end{aligned} \quad (6)$$

for $\mathbf{Q} = \mathbf{A}\mathbf{D}(\mathbf{s})^{-1}\mathbf{D}(\mathbf{x})\mathbf{A}^\top + \mathbf{D}(\mathbf{w})^{-1}\mathbf{D}(\mathbf{r})$. Thus, by solving the linear system in Equation (6), we can find $\Delta\mathbf{w}$, then $\Delta\mathbf{x}$, $\Delta\mathbf{r}$, and $\Delta\mathbf{s}$ through Equations (3)-(5).

The two algorithms we consider only differ in how they compute μ and α . The theoretical algorithm recomputes, at every iteration, $\mu = (\mathbf{x}^\top\mathbf{s} + \mathbf{w}^\top\mathbf{r})/(n+m)$, and chooses α to be the largest $\alpha < 1$ such that $x'_i s'_i \geq \gamma(\mathbf{x}^\top\mathbf{s}' + \mathbf{w}^\top\mathbf{r}')/(n+m)$ and $w'_j r'_j \geq \gamma(\mathbf{x}^\top\mathbf{s}' + \mathbf{w}^\top\mathbf{r}')/(n+m)$ for $i \in [n]$ and $j \in [m]$ for some hyperparameter $\gamma \in (0, 1]$ (Nocedal and Wright, 2006, Algorithm 14.2). The practical algorithm instead picks $\mu_0 = (\mathbf{x}_0^\top\mathbf{s}_0 + \mathbf{w}_0^\top\mathbf{r}_0)/(n+m)$ initially, for $(\mathbf{x}_0, \mathbf{w}_0, \mathbf{s}_0, \mathbf{r}_0)$ the initial point, and thereafter decreases μ as $\mu' = \sigma\mu$ at every iteration, while choosing α to be $\alpha = 0.99\alpha'$ for α' the largest $\alpha > 0$ such that $x'_i s'_i > 0$, $w'_j r'_j > 0$. The two algorithms are summarized in Algorithms 1 and 2.

Algorithm 1 is guaranteed to converge to an ϵ -accurate solution in $\mathcal{O}((n+m)\log(1/\epsilon))$ iterations (Nocedal and Wright, 2006, Theorem 14.3), that is, to a number of iterations proportional to the problem size. Algorithm 2, in contrast, does not come with any theoretical guarantees but is typical of practical IPM algorithms, which tend to converge in an almost constant number of

Algorithm 1 Theoretical IPM for LPs

Input: An LP instance $(\mathbf{A}, \mathbf{b}, \mathbf{c})$, a barrier reduction hyperparameter $\sigma \in (0, 1)$, a neighborhood hyperparameter $\gamma \in (0, 1]$, and initial values $(\mathbf{x}_0, \mathbf{w}_0, \mathbf{s}_0, \mathbf{r}_0)$ such that $\mathbf{A}\mathbf{x}_0 - \mathbf{r}_0 = \mathbf{b}$, $\mathbf{A}^\top\mathbf{w}_0 + \mathbf{s}_0 = \mathbf{c}$, $(\mathbf{x}_0, \mathbf{w}_0, \mathbf{s}_0, \mathbf{r}_0) > \mathbf{0}$ and $\min_i x_{0i} w_{0i} \geq \gamma\mu_0$, $\min_i w_{0i} r_{0i} \geq \gamma\mu_0$ for $\mu_0 = (\mathbf{x}_0^\top\mathbf{s}_0 + \mathbf{w}_0^\top\mathbf{r}_0)/(n+m)$.

- 1: **repeat**
- 2: $\mu \leftarrow (\mathbf{x}^\top\mathbf{s} + \mathbf{w}^\top\mathbf{r})/(n+m)$
- 3: Compute $\Delta\mathbf{w}$ by solving the linear system $\mathbf{Q}\Delta\mathbf{w} = \mathbf{b} - \mathbf{A}\mathbf{x} + \sigma\mu\mathbf{D}(\mathbf{w})^{-1}\mathbf{1} + \mathbf{A}\mathbf{D}(\mathbf{s})^{-1}\mathbf{D}(\mathbf{x})[\mathbf{c} - \mathbf{A}^\top\mathbf{w} - \sigma\mu\mathbf{D}(\mathbf{x})^{-1}\mathbf{1}]$ for $\mathbf{Q} = \mathbf{A}\mathbf{D}(\mathbf{s})^{-1}\mathbf{D}(\mathbf{x})\mathbf{A}^\top + \mathbf{D}(\mathbf{w})^{-1}\mathbf{D}(\mathbf{r})$
- 4: $\Delta\mathbf{x} \leftarrow \mathbf{D}(\mathbf{s})^{-1}\mathbf{D}(\mathbf{x})[\mathbf{A}^\top\Delta\mathbf{w} - \mathbf{c} + \mathbf{A}^\top\mathbf{w} + \sigma\mu\mathbf{D}(\mathbf{x})^{-1}\mathbf{1}]$
- 5: $\Delta\mathbf{s} \leftarrow \sigma\mu\mathbf{D}(\mathbf{x})^{-1}\mathbf{1} - \mathbf{s} - \mathbf{D}(\mathbf{x})^{-1}\mathbf{D}(\mathbf{s})\Delta\mathbf{x}$
- 6: $\Delta\mathbf{r} \leftarrow \sigma\mu\mathbf{D}(\mathbf{w})^{-1}\mathbf{1} - \mathbf{r} - \mathbf{D}(\mathbf{w})^{-1}\mathbf{D}(\mathbf{r})\Delta\mathbf{w}$
- 7: Compute the largest $\alpha \in (0, 1)$ such that

$$\begin{aligned} &\min_{i,j} \{(\mathbf{x} + \alpha\Delta\mathbf{x})_i (\mathbf{s} + \alpha\Delta\mathbf{s})_i, (\mathbf{w} + \alpha\Delta\mathbf{w})_j (\mathbf{r} + \alpha\Delta\mathbf{r})_j\} \\ &\geq \gamma \frac{(\mathbf{x} + \alpha\Delta\mathbf{x})^\top (\mathbf{s} + \alpha\Delta\mathbf{s}) + (\mathbf{w} + \alpha\Delta\mathbf{w})^\top (\mathbf{r} + \alpha\Delta\mathbf{r})}{n+m}. \end{aligned}$$

- 8: Update $(\mathbf{x}, \mathbf{w}, \mathbf{s}, \mathbf{r}) += \alpha(\Delta\mathbf{x}, \Delta\mathbf{w}, \Delta\mathbf{s}, \Delta\mathbf{r})$
 - 9: **until** convergence of $(\mathbf{x}, \mathbf{w}, \mathbf{s}, \mathbf{r})$
 - 10: **return** the point \mathbf{x} , which solves 1.
-

Algorithm 2 Practical IPM for LPs

Input: An LP instance $(\mathbf{A}, \mathbf{b}, \mathbf{c})$, a barrier reduction hyperparameter $\sigma \in (0, 1)$, and initial values $(\mathbf{x}_0, \mathbf{w}_0, \mathbf{s}_0, \mathbf{r}_0, \mu_0)$ such that $(\mathbf{x}_0, \mathbf{w}_0, \mathbf{s}_0, \mathbf{r}_0) > \mathbf{0}$ and $\mu_0 = (\mathbf{x}_0^\top\mathbf{s}_0 + \mathbf{w}_0^\top\mathbf{r}_0)/(n+m)$.

- 1: **repeat**
- 2: Compute $\Delta\mathbf{w}$ by solving the linear system $\mathbf{Q}\Delta\mathbf{w} = \mathbf{b} - \mathbf{A}\mathbf{x} + \sigma\mu\mathbf{D}(\mathbf{w})^{-1}\mathbf{1} + \mathbf{A}\mathbf{D}(\mathbf{s})^{-1}\mathbf{D}(\mathbf{x})[\mathbf{c} - \mathbf{A}^\top\mathbf{w} - \sigma\mu\mathbf{D}(\mathbf{x})^{-1}\mathbf{1}]$ for $\mathbf{Q} = \mathbf{A}\mathbf{D}(\mathbf{s})^{-1}\mathbf{D}(\mathbf{x})\mathbf{A}^\top + \mathbf{D}(\mathbf{w})^{-1}\mathbf{D}(\mathbf{r})$.
- 3: $\Delta\mathbf{x} \leftarrow \mathbf{D}(\mathbf{s})^{-1}\mathbf{D}(\mathbf{x})[\mathbf{A}^\top\Delta\mathbf{w} - \mathbf{c} + \mathbf{A}^\top\mathbf{w} + \sigma\mu\mathbf{D}(\mathbf{x})^{-1}\mathbf{1}]$
- 4: $\Delta\mathbf{s} \leftarrow \sigma\mu\mathbf{D}(\mathbf{x})^{-1}\mathbf{1} - \mathbf{s} - \mathbf{D}(\mathbf{x})^{-1}\mathbf{D}(\mathbf{s})\Delta\mathbf{x}$
- 5: $\Delta\mathbf{r} \leftarrow \sigma\mu\mathbf{D}(\mathbf{w})^{-1}\mathbf{1} - \mathbf{r} - \mathbf{D}(\mathbf{w})^{-1}\mathbf{D}(\mathbf{r})\Delta\mathbf{w}$
- 6: Find the largest $\alpha > 0$ such that

$$\begin{aligned} &\min_{i,j} \{(\mathbf{x} + \alpha\Delta\mathbf{x})_i (\mathbf{s} + \alpha\Delta\mathbf{s})_i, \\ &(\mathbf{w} + \alpha\Delta\mathbf{w})_j (\mathbf{r} + \alpha\Delta\mathbf{r})_j\} \geq 0 \end{aligned}$$

- 7: Update $(\mathbf{x}, \mathbf{w}, \mathbf{s}, \mathbf{r}) += 0.99\alpha(\Delta\mathbf{x}, \Delta\mathbf{w}, \Delta\mathbf{s}, \Delta\mathbf{r})$
 - 8: $\mu \leftarrow \sigma\mu$
 - 9: **until** convergence of $(\mathbf{x}, \mathbf{w}, \mathbf{s}, \mathbf{r})$
 - 10: **return** the point \mathbf{x} , which solves 1.
-

iterations—usually within 30-40 iterations, irrespective of problem size (Gondzio, 2012; Colombo and Gondzio, 2008).

Message-passing graph neural networks Intuitively, MPNNs learn node features or attributes, i.e., a d -element real-valued vector, representing each node in a graph by aggregating information from neighboring nodes. Let $\mathbf{G} = (G, \mathbf{L})$ be an attributed graph, following, Gilmer et al. (2017) and Scarselli et al. (2009), in each layer, $t > 0$, we update node attributes or features

$$\mathbf{h}_v^{(t)} := \text{UPD}^{(t)}\left(\mathbf{h}_v^{(t-1)}, \text{MSG}^{(t)}(\{\{\mathbf{h}_u^{(t-1)} \mid u \in N(v)\}\})\right),$$

and $\mathbf{h}_v^{(0)} := \mathbf{L}_v$. Here, the *message function* $\text{MSG}^{(t)}$ is a parameterized function, e.g., a neural network, mapping the multiset of neighboring node features to a single vectorial representation. We can easily adapt a message function to incorporate possible edge features or weights. Similarly, the *update function* $\text{UPD}^{(t)}$ is a parameterized function mapping the previous node features, and the output of $\text{MSG}^{(t)}$ to a single vectorial representation.

To adapt the parameters of the above functions, they are optimized end-to-end, usually through a variant of stochastic gradient descent, e.g., Kingma and Ba (2015), together with the parameters of a neural network used for classification or regression. In the following, we define a *message-passing step* as the application of a message and update function.

3 SIMULATING IPMS VIA MPNNs

We now show that there exist MPNNs, with specific architecture and choices of parameters, such that Algorithms 1 and 2 can be interpreted as inference over these MPNNs for a specific tripartite graph encoding the LP as input.

Representing LPs as graphs Let $I = (\mathbf{A}, \mathbf{b}, \mathbf{c})$ be an instance of LP. Similar to the setting of Ding et al. (2020), we model the instances with an (undirected) *weighted tripartite graph* $G(I) := (V(I), C(I), \{o\}, E(I)_{vc}, E(I)_{vo}, E(I)_{co})$. Here, the node set $V(I) := \{v_i \mid i \in [n]\}$ represents the variables of I , the node set $C(I) := \{c_i \mid i \in [m]\}$ represents the constraints of I , and the node o represents the objective. Further, the first edge set $E(I)_{vc}$ models the variable-constraint interaction, i.e., $E(I)_{vc} := \{(v_i, c_j) \mid A_{ij} \neq 0\}$, where each such edge (v_i, c_j) is annotated with the weight A_{ij} . Further, the objective node o is connected to all other nodes in the graphs, i.e., $E(I)_{vo} := \{(o, v_i) \mid v_i \in V(I)\}$ and $E(I)_{co} := \{(o, c_i) \mid c_i \in C(I)\}$. Each edge

$(o, v_i) \in E(I)_{vo}$ is annotated with the weight c_i . Similarly, each edge $(o, c_i) \in E(I)_{co}$ is annotated with the weight b_i . The resulting graph is illustrated in Figure 1a.

Theoretical results We now state the main results of this paper. To describe them, first notice that Algorithms 1 and 2 operate by taking an initial point $(\mathbf{x}_0, \mathbf{w}_0, \mathbf{s}_0, \mathbf{r}_0)$ and a duality measure μ_0 , and updating them after every iteration, yielding a sequence of points $(\mathbf{x}_t, \mathbf{w}_t, \mathbf{s}_t, \mathbf{r}_t)$ and duality measure μ_t for iterations $t > 0$. The following result shows that Algorithm 1 can be reproduced by a specific MPNN in the sense that a fixed-depth MPNN can reproduce each of its iterations.

Theorem 1. There exists an MPNN $f_{\text{MPNN,IPM1}}$ composed of $\mathcal{O}(m)$ message-passing steps that reproduces an iteration of Algorithm 1, in the sense that for any LP instance $I = (\mathbf{A}, \mathbf{b}, \mathbf{c})$ and any iteration step $t \geq 0$, $f_{\text{MPNN,IPM1}}$ maps the graph $G(I)$ carrying $[\mathbf{x}_t, \mathbf{s}_t]$ on the variable nodes and $[\mathbf{w}_t, \mathbf{r}_t]$ on the constraint nodes to the same graph $G(I)$ carrying $[\mathbf{x}_{t+1}, \mathbf{s}_{t+1}]$ on the variable nodes and $[\mathbf{w}_{t+1}, \mathbf{r}_{t+1}]$ on the constraint nodes.

This implies, by composing several instances of $f_{\text{MPNN,IPM1}}$, that Algorithm 1 can be simulated by an MPNN with a number of layers proportional to the number of iterations taken by the algorithm. We can derive an analogous result for Algorithm 2.

Proposition 2. There exists an MPNN $f_{\text{MPNN,IPM2}}$ composed of $\mathcal{O}(m)$ message-passing steps that reproduces each iteration of Algorithm 2, in the sense that for any LP instance $I = (\mathbf{A}, \mathbf{b}, \mathbf{c})$ and any iteration step $t \geq 0$, $f_{\text{MPNN,IPM2}}$ maps the graph $G(I)$ carrying $[\mathbf{x}_t, \mathbf{s}_t]$ on the variable nodes, $[\mathbf{w}_t, \mathbf{r}_t]$ on the constraint nodes and $[\mu_t]$ on the objective node to the same graph $G(I)$ carrying $[\mathbf{x}_{t+1}, \mathbf{s}_{t+1}]$ on the variable nodes, $[\mathbf{w}_{t+1}, \mathbf{r}_{t+1}]$ on the constraint nodes and $[\mu_{t+1}]$ on the objective node.

Similarly, this implies, by composing several instances of $f_{\text{MPNN,IPM2}}$, that Algorithm 2 can be simulated by an MPNN with a number of layers proportional to the number of iterations taken by the algorithm.

Implication of the theoretical results Theorem 1 and Proposition 2 show that MPNNs are, in principle, capable of simulating modern IPMs. That is, they are capable of solving LPs to optimality. Hence, our findings shed light on the recent success of MPNN-based neural architectures by Gasse et al. (2019) and similar approaches, which use MPNNs to mimic strong branching within the branch-and-bound framework for solving mixed-integer linear optimization problems. Moreover, in the following section, we derive MPNN architectures that act as lightweight proxies for solving LPs while

being able to adapt to the given problem instance distributions. We verify their effectiveness empirically on real-world LP instances stemming from relaxing mixed-integer linear formulations of well-known combinatorial optimization problems.

4 IPM-MPNNs: MPNNs FOR LPS

Inspired by the theoretical alignment of IPMs and MPNNs derived above, we outline our IPM-MPNN framework, allowing for solving LP instances while adapting to a given problem instance distribution.

Given an LP instance I , we now outline an asynchronous MPNN operating on the tripartite graph $G(I)$; see Section 3. Let $\mathbf{h}_c^{(t)} \in \mathbb{R}^d$, $d > 0$, be the node features of a constraint node $c \in C(I)$ at iteration $t > 0$, and let $\mathbf{h}_v^{(t)} \in \mathbb{R}^d$ and $\mathbf{h}_o^{(t)} \in \mathbb{R}^d$ be the node features of a variable node $v \in V(I)$ and the objective node o at iteration t , respectively. Moreover, let $\mathbf{e}_{co}, \mathbf{e}_{vc}, \mathbf{e}_{vo}$ denote the edge weights. Initially, at $t = 0$, we set the node features by applying a linear mapping to the raw node features $\mathbf{x}_v, \mathbf{x}_c$ or \mathbf{x}_o , extracted from the instance I ; see Section 5 for details.

All three node types are updated in three separate update passes. In the first pass, we update the embeddings of constraint nodes from the embeddings of the variable nodes and of the objective node. That is, let $c \in C(I)$ be a constraint node and let $t > 0$, then

$$\mathbf{h}_c^{(t)} := \text{UPD}_c^{(t)} \left[\mathbf{h}_c^{(t-1)}, \text{MSG}_{o \rightarrow c}^{(t)} \left(\mathbf{h}_o^{(t-1)}, \mathbf{e}_{oc} \right), \text{MSG}_{v \rightarrow c}^{(t)} \left(\{ \mathbf{h}_v^{(t-1)}, \mathbf{e}_{vc} \} \mid v \in N(c) \cap V(I) \} \right) \right].$$

Here, the parameterized message function $\text{MSG}_{v \rightarrow c}^{(t)}$ maps a multiset of vectors, i.e., variable node features and corresponding edge features \mathbf{e}_{vc} , to a vector in \mathbb{R}^d . Similarly, the parameterized function $\text{MSG}_{o \rightarrow c}$ maps the current node features of the objective node and edge features \mathbf{e}_{oc} to a vector in \mathbb{R}^d . Finally, the parameterized function $\text{UPD}_c^{(t)}$ maps the constraint node's previous features, the outputs of $\text{MSG}_{o \rightarrow c}^{(t)}$ and $\text{MSG}_{v \rightarrow c}^{(t)}$ to vector in \mathbb{R}^d .

Next, similarly to the above, we update the objective node's features depending on variable and constraint node features,

$$\mathbf{h}_o^{(t)} := \text{UPD}_o^{(t)} \left[\mathbf{h}_o^{(t-1)}, \text{MSG}_{c \rightarrow o}^{(t)} \left(\{ \mathbf{h}_c^{(t)}, \mathbf{e}_{co} \} \mid c \in C(I) \} \right), \text{MSG}_{v \rightarrow o}^{(t)} \left(\{ \mathbf{h}_v^{(t-1)}, \mathbf{e}_{vo} \} \mid v \in V(I) \} \right) \right].$$

Finally, analogously to the update of the constraint nodes' features, we update the representation of a variable node $v \in V(I)$ from the constraints nodes and

objective node,

$$\mathbf{h}_v^{(t)} := \text{UPD}_v^{(t)} \left[\mathbf{h}_v^{(t-1)}, \text{MSG}_{o \rightarrow v}^{(t)} \left(\mathbf{h}_o^{(t)}, \mathbf{e}_{ov} \right), \text{MSG}_{c \rightarrow v}^{(t)} \left(\{ \mathbf{h}_c^{(t)}, \mathbf{e}_{cv} \} \mid c \in N(v) \cap V(C) \} \right) \right].$$

The entire process is executed asynchronously, in that the nodes updated later incorporate the most recent features updates from preceding nodes. We map each variable node feature $\mathbf{h}_v^{(t)}$ to $\text{MLP}(\mathbf{h}_v^{(t)}) \in \mathbb{R}$, where MLP is a multi-layer perceptron, and concatenate the resulting real numbers over all variable nodes column-wise, resulting in the final prediction $\mathbf{z}^{(t)} \in \mathbb{R}^n$.

In the experiments in Section 5, we probe various message-passing layers to express the various message and update functions. Concretely, we leverage the GCN (Kipf and Welling, 2017), GIN (Xu et al., 2019), and GEN (Li et al., 2020) MPNN layers, respectively; see Appendix B for details.

We train the above IPM-MPNN architecture, i.e., adapt its parameters, in a supervised fashion. Below, we outline the three components constituting our training loss function. To that, let $(\mathbf{A}, \mathbf{b}, \mathbf{c})$ be an LP instance.

Variable supervision As discussed above, our IPM-MPNN aims to simulate the solving steps provided by a standard IPM. Thus, aligned with Theorem 1 and Proposition 2, a perfectly parameterized MPNN is expected to follow the steps without deviation for all T iterations. We maintain the intermediate outputs $\mathbf{z}^{(t)} \in \mathbb{R}^n$ of each MPNN layer and calculate the mean squared error (MSE) loss between every pair of the expert solution $\mathbf{y}^{(t)} \in \mathbb{R}^n$ and MPNN prediction $\mathbf{z}^{(t)}$. Moreover, we introduce a step decay factor $\alpha \in [0, 1]$ so that early steps play a less important role, resulting in the following loss function,

$$\mathcal{L}_{\text{var}} := \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T \alpha^{T-t} \|\mathbf{y}_i^{(t)} - \mathbf{z}_i^{(t)}\|_2^2, \quad (7)$$

where N denotes the number of training samples.

Objective supervision We use regularization on the prediction regarding the ground-truth objective values at every step. Empirically, this regularization term helps with convergence and helps finding more feasible solutions that minimize the objective in case the LP instance has multiple solutions. Note that we do not predict the objective directly with our MPNNs but calculate it via $\mathbf{c}^\top \mathbf{z}^{(t)}$ instead. Suppose the ground-truth values are given by $\mathbf{c}^\top \mathbf{y}^{(t)}$, we have

$$\mathcal{L}_{\text{obj}} := \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T \alpha^{T-t} \left[\mathbf{c}^\top (\mathbf{y}_i^{(t)} - \mathbf{z}_i^{(t)}) \right]^2.$$

Constraint supervision Finally, we aim for the IPM-MPNN to predict an optimal solution regarding the objective value while satisfying all the constraints. To that, we introduce a regularization penalizing constraint violations, i.e.,

$$\mathcal{L}_{\text{cons}} := \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T \alpha^{T-t} \|\text{ReLU}(\mathbf{A}_i \mathbf{z}_i^{(t)} - \mathbf{b}_i)\|_2^2.$$

Finally, we combine the above three loss terms into the loss function

$$\mathcal{L} := w_{\text{var}} \mathcal{L}_{\text{var}} + w_{\text{obj}} \mathcal{L}_{\text{obj}} + w_{\text{cons}} \mathcal{L}_{\text{cons}},$$

where we treat $w_{\text{var}}, w_{\text{obj}}$, and $w_{\text{cons}} > 0$ as hyperparameters.

At test time, given an LP instance I , we construct the tripartite graph $G(I)$ as outlined above and use the trained MPNN to predict the variables’ values.

5 EXPERIMENTAL STUDY

Here, we empirically evaluate the ability of IPM-MPNNs to predict the optimal solutions of LPs. In particular, we aim to answer the following questions.

- Q1** Can MPNNs properly imitate the performance of IPM solvers in practice?
- Q2** How is MPNNs’ performance compared with competing neural-network-based solutions?
- Q3** What advantage does our MPNN architecture hold compared with traditional IPM solvers?
- Q4** Do MPNNs generalize to instances larger than seen during training?

Our experimental results are reproducible with the code available at https://github.com/chendiqian/IPM_MPNN.

Datasets We obtain LP instances from mixed-integer optimization instances by dropping the integrality constraints over variables. Following Gasse et al. (2019), we use four classes of problems, namely *set covering*, *maximal independent set*, *combinatorial auction*, and *capacitated facility location*. For each problem type, we generate small- and large-size instances. We describe the datasets and our parameters for dataset generation in Appendix C. From each LP instance $I = (\mathbf{A}, \mathbf{b}, \mathbf{c})$, we generate a tripartite graph $G(I)$, following Section 3, and construct initial node and edge features as follows. We indicate the i th row of the constraint matrix \mathbf{A} as \mathbf{A}_i and the j th column as $\mathbf{A}_{\cdot,j}$. For a given variable node $v_j \in V(I)$, its initial node features are set to the mean and standard deviation of the column vector $\mathbf{A}_{\cdot,j}$,

resulting in two features. Analogously, for a constraint node $c_i \in C(I)$, we derive initial features from the statistical properties of the row vector \mathbf{A}_i . For the objective node o , the features are characterized in a corresponding manner using the vector \mathbf{c} . In the supervised learning regime, it is mandatory to have ground truth labels to guide the model predictions. For our IPM-MPNNs, the outputs of each layer t are the prediction of the variables’ value $\mathbf{z}^{(t)}$. Consequently, we utilize the intermediate variable values $\mathbf{y}^{(t)}$, as provided by the solver, to serve as our ground truth. However, to prevent the MPNN from becoming excessively deep, we sample the ground truth steps, i.e., we adopt an equidistant sampling strategy for the solver’s steps, ensuring that the number of sampled steps aligns with the depth of the corresponding MPNN. We split the graph datasets into train, validation, and test splits with ratios 0.8, 0.1, and 0.1. We conducted the experiments by evaluating each dataset multiple times, performing three independent runs with distinct random seeds. The reported results represent the average numbers over the test set across these runs and the corresponding standard deviations. We executed all experiments on a single NVIDIA A100 80GB GPU card—see Appendix D for training-related hyperparameters.

IPM-MPNNs’ ability to solve LPs (Q1) We collect and organize the results of our MPNN method in Table 1. We report the numbers on the four types of relaxed MILP instances, each with small and large sizes. As explained in Section 4, we leverage three types of MPNN layers, GCN (Kipf and Welling, 2017), GIN (Xu et al., 2019), and GEN (Li et al., 2020), as the backbone of our MPNN architectures. We split the table into two main parts, namely the mean absolute relative objective gap

$$\frac{1}{N} \sum_{i=1}^N \left| \frac{\mathbf{c}^\top (\mathbf{y}_i^{(T)} - \mathbf{z}_i^{(T)})}{\mathbf{c}^\top \mathbf{y}_i^{(T)}} \right| \times 100$$

of the last step T over the test set, and the mean absolute constraint violation of the last step

$$\frac{1}{N} \sum_{i=1}^N \frac{1}{m_i} \|\text{ReLU}(\mathbf{A}_i \mathbf{z}_i^{(T)} - \mathbf{b}_i)\|_1,$$

where the normalization term m_i is the number of constraints of the i th instance. As seen, our IPM-MPNN architectures consistently align with the IPM solver at the last converged step, with marginal constraint violation. Moreover, the relative objective gaps of our method are all under 1%. The GCN-based IPM-MPNNs perform best on large maximal independent set relaxation instances at $0.094 \pm 0.005\%$. Overall, GEN and GCN perform better than GIN layers among the three graph convolutions. The absolute constraint

Table 1: Results of our proposed IPM-MPNNs (\checkmark) versus bipartite representation ablations (\times). We report the relative objective gap and the constraint violation, averaged over all three runs. We print the best results per target in bold.

| | Tri. | MPNN | Small instances | | | Large instances | | | | |
|----------------------|--------------|------|------------------------------------|-------------------------------------|------------------------------------|------------------------------------|-----------------------------------|--------------------------------------|-------------------------------------|------------------------------------|
| | | | Setcover | Indset | Cauc | Fac | Setcover | Indset | Cauc | Fac |
| Objective gap [%] | \checkmark | GEN | 0.319\pm0.020 | 0.119 \pm 0.003 | 0.612\pm0.049 | 0.549\pm0.112 | 0.629 \pm 0.086 | 0.158 \pm 0.035 | 0.306\pm0.047 | 0.747\pm0.083 |
| | | GCN | 0.418 \pm 0.008 | 0.103\pm0.006 | 0.682 \pm 0.029 | 0.578 \pm 0.015 | 0.420\pm0.047 | 0.094\pm0.005 | 0.407 \pm 0.038 | 0.914 \pm 0.141 |
| | | GIN | 0.478 \pm 0.038 | 0.146 \pm 0.011 | 0.632 \pm 0.036 | 0.810 \pm 0.221 | 0.711 \pm 0.115 | 0.126 \pm 0.021 | 0.378 \pm 0.052 | 0.911 \pm 0.132 |
| | \times | GEN | 8.310 \pm 1.269 | 0.735 \pm 0.032 | 1.417 \pm 0.009 | 2.976 \pm 0.013 | 15.170 \pm 6.844 | 0.320 \pm 0.008 | 0.851 \pm 0.122 | 2.531 \pm 0.025 |
| | | GCN | 5.523 \pm 0.133 | 0.639 \pm 0.009 | 1.394 \pm 0.081 | 3.031 \pm 0.059 | 6.092 \pm 0.456 | 0.298 \pm 0.009 | 0.766 \pm 0.093 | 2.535 \pm 0.034 |
| | | GIN | 5.592 \pm 0.179 | 0.634 \pm 0.021 | 1.202 \pm 0.016 | 2.996 \pm 0.031 | 5.835 \pm 1.917 | 0.290 \pm 0.005 | 0.810 \pm 0.140 | 2.660 \pm 0.062 |
| Constraint violation | \checkmark | GEN | 0.002\pm0.0002 | 0.0006 \pm 0.00003 | 0.003 \pm 0.0007 | 0.002 \pm 0.001 | 0.009 \pm 0.001 | 0.0015 \pm 0.0003 | 0.0004\pm0.0002 | 0.002 \pm 0.001 |
| | | GCN | 0.002 \pm 0.001 | 0.0003\pm0.0001 | 0.002 \pm 0.00007 | 0.002\pm0.0002 | 0.009 \pm 0.001 | 0.0005\pm0.00004 | 0.001 \pm 0.0005 | 0.001\pm0.0004 |
| | | GIN | 0.004 \pm 0.001 | 0.0006 \pm 0.00008 | 0.001\pm0.0001 | 0.002 \pm 0.0005 | 0.008\pm0.002 | 0.0006 \pm 0.0001 | 0.002 \pm 0.0008 | 0.002 \pm 0.0007 |
| | \times | GEN | 0.181 \pm 0.023 | 0.006 \pm 0.0003 | 0.006 \pm 0.001 | 0.011 \pm 0.004 | 0.309 \pm 0.025 | 0.004 \pm 0.0002 | 0.006 \pm 0.001 | 0.003 \pm 0.001 |
| | | GCN | 0.207 \pm 0.006 | 0.004 \pm 0.001 | 0.002 \pm 0.001 | 0.006 \pm 0.0003 | 0.267 \pm 0.049 | 0.003 \pm 0.0004 | 0.004 \pm 0.001 | 0.002 \pm 0.0003 |
| | | GIN | 0.211 \pm 0.007 | 0.003 \pm 0.0002 | 0.003 \pm 0.001 | 0.008 \pm 0.002 | 0.236 \pm 0.014 | 0.003 \pm 0.0004 | 0.004 \pm 0.002 | 0.003 \pm 0.0002 |

violations are mostly at the 1×10^{-3} level, with the best (0.0003 ± 0.0001) achieved by GCN on small maximal independent set relaxation instances.

Baselines (Q2) To answer question Q2, we compare IPM-MPNNs to two baselines. First, we compare our IPM-MPNNs to Chen et al. (2023), where the authors proposed encoding LP instances as bipartite graphs. In Table 1, the rows with a cross mark (\times) on the left are the bipartite baselines. As seen in this table, our IPM-MPNNs outperform the bipartite architectures in all instances with all types of MPNNs. Most relative objective gaps lie above 1%, except on the maximal independent set and large combinatorial auction relaxation instances. The largest gap is observed on the small set covering relaxation instances by GEN, with the baseline reporting as much as $26.1 \times$ higher constraint violation number. Hence, our results indicate that IPM-MPNNs’s tripartite representation is crucial. We also compare our IPM-MPNNs to a neural-ODE-based approach (Wu and Lissner, 2023). Since their approach is quite expensive during training by means of both time and GPU memory, we generate 1000 mini-sized instances. Due to the architecture-agnostic property of their approach, we embed our MPNNs in their training pipeline. When presenting the runtime and GPU usage, we use the same batch size as the baseline method on our MPNN approach for a fair comparison, even though our method can scale to a much larger batch size in practice. We report results in Table 2. Taking the GEN layer as an example, our method shows consistently better results than the neural-ODE baseline method by reaching at most $14.2 \times$ lower relative objective gap, $8.6 \times$ faster training, and $300.8 \times$ less memory on the mini-sized capacitated facility location relaxation instances.

Table 2: Comparing between IPM-MPNNs and the Wu and Lissner (2023) method on 1000 mini-sized instances. We report the average relative objective gap, constraint violation, training time over three runs, and maximal GPU memory allocated. We print the best results per target in bold.

| | Method | MPNN | Setcover | Indset | Cauc | Fac |
|-----------------|--------|-----------------------------------|-----------------------------------|------------------------------------|-----------------------------------|------------------------------------|
| Obj. gap [%] | ODE | GEN | 14.915 \pm 0.425 | 6.225 \pm 0.097 | 13.845 \pm 0.554 | 20.560 \pm 0.059 |
| | | GCN | 14.545 \pm 0.055 | 6.148 \pm 0.071 | 12.945 \pm 0.385 | 20.690 \pm 0.037 |
| | | GIN | 15.050 \pm 0.228 | 6.474 \pm 0.114 | 13.470 \pm 0.145 | 21.010 \pm 0.529 |
| Ours | GEN | 2.555 \pm 0.122 | 1.580 \pm 0.095 | 2.733\pm0.074 | 1.449 \pm 0.255 | |
| | GCN | 2.375\pm0.062 | 1.447 \pm 0.152 | 2.769 \pm 0.091 | 1.478 \pm 0.154 | |
| | GIN | 2.740 \pm 0.3184 | 1.404\pm0.153 | 2.847 \pm 0.091 | 1.328\pm0.201 | |
| Constraint vio. | ODE | GEN | 0.072 \pm 0.006 | 0.046 \pm 0.002 | 0.025 \pm 0.008 | 0.020 \pm 0.001 |
| | | GCN | 0.049 \pm 0.012 | 0.048 \pm 0.008 | 0.025 \pm 0.0002 | 0.020 \pm 0.0005 |
| | | GIN | 0.064 \pm 0.005 | 0.043 \pm 0.008 | 0.024 \pm 0.005 | 0.014 \pm 0.004 |
| | Ours | GEN | 0.023\pm0.002 | 0.005 \pm 0.0001 | 0.015 \pm 0.003 | 0.013 \pm 0.003 |
| | | GCN | 0.030 \pm 0.003 | 0.005 \pm 0.0006 | 0.017 \pm 0.002 | 0.005\pm0.0006 |
| | | GIN | 0.023 \pm 0.005 | 0.005\pm0.0003 | 0.014\pm0.001 | 0.006 \pm 0.0006 |
| Time [s] | ODE | GEN | 47.829 | 51.283 | 63.068 | 96.298 |
| | | GCN | 57.196 | 80.133 | 79.606 | 34.297 |
| | | GIN | 55.918 | 64.628 | 39.904 | 62.448 |
| | Ours | GEN | 10.177 | 9.617 | 9.946 | 11.124 |
| | | GCN | 18.964 | 8.688 | 7.368 | 8.834 |
| | | GIN | 6.042 | 8.096 | 8.881 | 10.771 |
| Memory (GB) | ODE | GEN | 16.455 | 25.931 | 23.354 | 44.520 |
| | | GCN | 16.489 | 34.003 | 23.805 | 10.640 |
| | | GIN | 18.238 | 30.101 | 13.482 | 24.713 |
| | Ours | GEN | 0.091 | 0.088 | 0.101 | 0.148 |
| | | GCN | 0.201 | 0.134 | 0.069 | 0.142 |
| | | GIN | 0.094 | 0.073 | 0.148 | 0.187 |

Inference time profiling (Q3) We also compare IPM-MPNN’s performance to exact IPM solvers. Thereto, we compare the time required to solve an instance between traditional solvers, namely SciPy’s IPM solver and a Python-based custom-build one, and our IPM-MPNN. We run the solvers and our MPNNs on the test set of each dataset and report the mean and standard deviation in seconds. According to Table 4, our MPNN clearly outperforms the SciPy IPM solver on all large instances. Further, IPM-MPNNs beat our Python-based IPM solver described in Algorithm 2 on all instances. It is worth noting that both IPM solvers exhibit sensitivity to problem sizes. For example, the SciPy solver sees a performance degradation of approx-

Table 3: Size generalization. We report the relative objective gap and constraint violation on larger test instances. Numbers represent mean and standard deviation across multiple pretrained models.

| | Train size | | Inference size | | GEN | | GCN | | GIN | |
|---------|------------|------------|----------------|------|-------------|-------------|-------------|-------------|---------------|-------------|
| | Rows | Cols | Rows | Cols | Obj. (%) | Cons. | Obj. (%) | Cons. | Obj. (%) | Cons. |
| Setc. | [300, 500] | [500, 700] | 500 | 700 | 0.717±0.158 | 0.516±0.010 | 0.511±0.047 | 0.509±0.004 | 1.034±0.237 | 0.486±0.023 |
| | | | 550 | 750 | 0.917±0.317 | 0.552±0.012 | 0.871±0.252 | 0.543±0.003 | 2.318±1.411 | 0.497±0.032 |
| | | | 600 | 700 | 0.993±0.211 | 0.573±0.015 | 0.705±0.125 | 0.565±0.012 | 1.491±0.512 | 0.521±0.045 |
| | | | 500 | 800 | 0.902±0.323 | 0.528±0.008 | 1.058±0.441 | 0.509±0.004 | 12.538±16.027 | 0.485±0.050 |
| | | | 600 | 800 | 1.004±0.407 | 0.589±0.014 | 1.556±0.588 | 0.568±0.005 | 12.217±14.715 | 0.486±0.071 |
| Indset. | [584, 990] | [300, 500] | [978, 994] | 500 | 0.128±0.027 | 0.299±0.001 | 0.099±0.008 | 0.303±0.001 | 0.129±0.031 | 0.304±0.001 |
| | | | [1028, 1044] | 525 | 0.157±0.063 | 0.300±0.001 | 0.101±0.013 | 0.304±0.001 | 0.111±0.017 | 0.305±0.001 |
| | | | [1076, 1094] | 550 | 0.300±0.186 | 0.301±0.002 | 0.096±0.022 | 0.303±0.001 | 0.177±0.097 | 0.304±0.001 |
| | | | [1128, 1144] | 575 | 1.402±1.036 | 0.305±0.006 | 0.146±0.044 | 0.304±0.001 | 0.380±0.367 | 0.304±0.002 |
| | | | [1178, 1194] | 600 | 4.552±3.153 | 0.317±0.015 | 0.408±0.317 | 0.304±0.001 | 0.647±0.725 | 0.304±0.002 |
| Cauc. | [320, 562] | [300, 499] | [530, 564] | 500 | 0.333±0.134 | 0.257±0.001 | 0.318±0.048 | 0.259±0.001 | 0.344±0.108 | 0.259±0.001 |
| | | | [596, 646] | 500 | 0.363±0.131 | 0.267±0.002 | 0.519±0.069 | 0.270±0.003 | 0.576±0.165 | 0.271±0.001 |
| | | | [652, 720] | 500 | 0.524±0.039 | 0.284±0.001 | 1.255±0.523 | 0.289±0.007 | 0.944±0.114 | 0.289±0.001 |
| | | | [559, 596] | 600 | 7.325±3.615 | 0.257±0.002 | 0.587±0.268 | 0.255±0.001 | 1.014±0.845 | 0.263±0.006 |
| | | | [633, 677] | 600 | 7.965±3.941 | 0.263±0.002 | 0.868±0.441 | 0.258±0.003 | 1.375±0.693 | 0.269±0.005 |
| Fac. | [441, 900] | [420, 870] | 961 | 930 | 0.912±0.251 | 0.178±0.006 | 1.154±0.206 | 0.173±0.007 | 1.452±0.528 | 0.178±0.003 |
| | | | 936 | 900 | 1.320±0.347 | 0.148±0.009 | 1.615±0.322 | 0.145±0.009 | 1.736±0.558 | 0.153±0.004 |
| | | | 936 | 910 | 0.964±0.063 | 0.209±0.005 | 1.538±0.526 | 0.211±0.007 | 1.538±0.422 | 0.215±0.006 |
| | | | 1116 | 1080 | 1.502±0.704 | 0.163±0.009 | 3.540±3.134 | 0.161±0.006 | 2.288±0.659 | 0.167±0.005 |
| | | | 1296 | 1260 | 1.808±0.566 | 0.173±0.009 | 7.629±7.577 | 0.179±0.008 | 13.522±8.027 | 0.163±0.021 |

Table 4: Comparing IPM-MPNNs’ inference time to SciPy’s IPM implementation and our Python-based IPM solver. We report mean and standard deviation in seconds over three runs. We print the best results per target in bold.

| Instances | SciPy Solver | Our Solver | GEN | GCN | GIN |
|----------------|--------------------|-------------|-------------|-------------|--------------------|
| Small setcover | 0.006±0.004 | 0.071±0.015 | 0.033±0.001 | 0.029±0.001 | 0.017±0.001 |
| Large setcover | 0.390±0.098 | 3.696±2.141 | 0.033±0.001 | 0.030±0.001 | 0.021±0.001 |
| Small indset | 0.008±0.007 | 0.089±0.024 | 0.033±0.001 | 0.031±0.002 | 0.021±0.001 |
| Large indset | 0.226±0.087 | 1.053±0.281 | 0.033±0.002 | 0.030±0.001 | 0.021±0.001 |
| Small cauc | 0.012±0.005 | 0.151±0.035 | 0.033±0.001 | 0.028±0.001 | 0.021±0.001 |
| Large cauc | 0.282±0.065 | 3.148±0.880 | 0.033±0.001 | 0.029±0.001 | 0.021±0.001 |
| Small fac | 0.017±0.011 | 2.025±1.854 | 0.029±0.001 | 0.029±0.001 | 0.022±0.001 |
| Large fac | 0.732±0.324 | 6.229±2.672 | 0.030±0.001 | 0.031±0.001 | 0.022±0.001 |

imately $65.0\times$ when transitioning from small to large set covering problems. In contrast, our MPNN method demonstrates a more consistent behavior across varying problem sizes, showing only $1.2\times$ slowdown for the analogous instances. Hence, we can positively answer question **Q3**.

Size generalization (Q4) We investigate the possibility of generalizing our pre-trained MPNNs to larger instances than encountered during training. To that, we generate new sets of novel instances, each with the same number of instances as the test set. In Table 3, we list our training instance sizes and a collection of test instance sizes. Taking GEN on set covering relaxation problems as an example, as the inference size grows, the relative objective gap increases a bit, while the constraint violations are overall worse than on the training set. Notably, for the training size [600,800], which is at least $1.37\times$ larger than the training instances, the objective gap is merely 0.28% worse than [500,700] size. For question **Q4**, we can conclude that our pre-trained

MPNNs can generalize to unseen, larger instances to some extent.

6 CONCLUSION

In summary, our study establishes a strong connection between MPNNs and IPMs for LPs. We have shown that MPNNs can effectively simulate IPM iterations, revealing their potential in emulating strong branching within the branch-and-bound framework, as demonstrated by Gasse et al. (2019). In addition, building on this connection, we proposed IPM-MPNNs for learning to solve large-scale LP instances approximately, surpassing neural baselines and exact IPM solvers in terms of solving time across various problem domains. Looking forward, promising avenues for further research involve expanding our theoretical framework to encompass a broader range of convex optimization problems.

Acknowledgments

CQ and CM are partially funded by a DFG Emmy Noether grant (468502433) and RWTH Junior Principal Investigator Fellowship under Germany’s Excellence Strategy.

References

- Achterberg, T., Koch, T., and Martin, A. (2005). Branching rules revisited. *Operations Research Letters*, 33(1):42–54.
- Baskin, I. I., Palyulin, V. A., and Zefirov, N. S. (1997). A neural device for searching direct correlations between structures and properties of chemical com-

- pounds. *Journal of Chemical Information and Computer Sciences*, 37(4):715–721.
- Bello, I., Pham, H., Le, Q. V., Norouzi, M., and Bengio, S. (2016). Neural combinatorial optimization with reinforcement learning. *CoRR*, abs/1611.09940.
- Bengio, Y., Frejinger, E., Lodi, A., Patel, R., and Sankaranarayanan, S. (2020). A learning-based algorithm to quickly compute good primal solutions for stochastic integer programs. In *International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, pages 99–111.
- Bengio, Y., Lodi, A., and Prouvost, A. (2021). Machine learning for combinatorial optimization: A methodological tour d’horizon. *Eur. J. Oper. Res.*, 290(2):405–421.
- Bruna, J., Zaremba, W., Szlam, A., and LeCun, Y. (2014). Spectral networks and deep locally connected networks on graphs. In *International Conference on Learning Representation*.
- Cappart, Q., Chételat, D., Khalil, E. B., Lodi, A., Morris, C., and Veličković, P. (2021). Combinatorial optimization and reasoning with graph neural networks. In *Joint Conference on Artificial Intelligence*, pages 4348–4355.
- Chen, Z., Liu, J., Wang, X., and Yin, W. (2023). On representing linear programs by graph neural networks. In *International Conference on Learning Representations*.
- Chmiela, A., Khalil, E., Gleixner, A., Lodi, A., and Pokutta, S. (2021). Learning to schedule heuristics in branch and bound. *Advances in Neural Information Processing Systems*, 34.
- Colombo, M. and Gondzio, J. (2008). Further development of multiple centrality correctors for interior point methods. *Computational Optimization and Applications*, 41:277–305.
- Cybenko, G. (1992). Approximation by superpositions of a sigmoidal function. *Math. Control. Signals Syst.*, 5(4):455.
- Defferrard, M., Bresson, X., and Vandergheynst, P. (2016). Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, pages 3837–3845.
- Deza, A. and Khalil, E. B. (2023). Machine learning for cutting planes in integer programming: A survey. In *Joint Conference on Artificial Intelligence*, pages 6592–6600.
- Ding, J.-Y., Zhang, C., Shen, L., Li, S., Wang, B., Xu, Y., and Song, L. (2020). Accelerating primal solution findings for mixed integer programs based on solution prediction. In *AAAI Conference on Artificial Intelligence*.
- Duvenaud, D., Maclaurin, D., Aguilera-Iparraguirre, J., Gómez-Bombarelli, R., Hirzel, T., Aspuru-Guzik, A., and Adams, R. P. (2015). Convolutional networks on graphs for learning molecular fingerprints. In *Advances in Neural Information Processing Systems*, pages 2224–2232.
- Fan, Z., Wang, X., Yakovenko, O., Sivas, A. A., Ren, O., Zhang, Y., and Zhou, Z. (2023). Smart initial basis selection for linear programs. In *International Conference on Machine Learning*, volume 202, pages 9650–9664.
- Gasse, M., Chételat, D., Ferroni, N., Charlin, L., and Lodi, A. (2019). Exact combinatorial optimization with graph convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 15554–15566.
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. (2017). Neural message passing for quantum chemistry. In *International Conference on Machine Learning*, pages 1263–1272.
- Gondzio, J. (2012). Interior point methods 25 years later. *European Journal of Operational Research*, 218:587–601.
- Gupta, P., Gasse, M., Khalil, E. B., Mudigonda, P. K., Lodi, A., and Bengio, Y. (2020). Hybrid models for learning to branch. In *Advances in Neural Information Processing Systems* 33.
- Gupta, P., Khalil, E. B., Chételat, D., Gasse, M., Lodi, A., Bengio, Y., and Kumar, M. P. (2022). Lookback for learning to branch. *Transactions on Machine Learning Research*, 2022.
- Hamilton, W. L., Ying, Z., and Leskovec, J. (2017). Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, pages 1024–1034.
- Khalil, E. B., Bodic, P. L., Song, L., Nemhauser, G. L., and Dilkina, B. (2016). Learning to branch in mixed integer programming. In *AAAI Conference on Artificial Intelligence*, pages 724–731.
- Khalil, E. B., Dilkina, B., Nemhauser, G. L., Ahmed, S., and Shao, Y. (2017). Learning to run heuristics in tree search. In *International Joint Conference on Artificial Intelligence*, pages 659–666.
- Khalil, E. B., Morris, C., and Lodi, A. (2022). MIP-GNN: A data-driven framework for guiding combinatorial solvers. In *AAAI Conference on Artificial Intelligence*, pages 10219–10227.

- Kingma, D. P. and Ba, J. (2015). Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- Kipf, T. N. and Welling, M. (2017). Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*.
- Kireev, D. B. (1995). Chemnet: A novel neural network based method for graph/property mapping. *Journal of Chemical Information and Computer Sciences*, 35(2):175–180.
- Kotary, J., Fioretto, F., Van Hentenryck, P., and Wilder, B. (2021). End-to-end constrained optimization learning: A survey. *CoRR*, abs/2103.16378.
- Leshno, M., Lin, V. Y., Pinkus, A., and Schocken, S. (1993). Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural Networks*, 6(6):861–867.
- Li, G., Xiong, C., Thabet, A., and Ghanem, B. (2020). Deepergcn: All you need to train deeper GCNs. *CoRR*, abs/2006.07739.
- Li, X., Qu, Q., Zhu, F., Zeng, J., Yuan, M., Mao, K., and Wang, J. (2022). Learning to reformulate for linear programming. *CoRR*, abs/2201.06216.
- Merkwirth, C. and Lengauer, T. (2005). Automatic generation of complementary descriptors with molecular graph networks. *Journal of Chemical Information and Modeling*, 45(5):1159–1168.
- Micheli, A. (2009). Neural network for graphs: A contextual constructive approach. *IEEE Transactions on Neural Networks*, 20(3):498–511.
- Micheli, A. and Sestito, A. S. (2005). A new neural network model for contextual processing of graphs. In *Italian Workshop on Neural Nets Neural Nets and International Workshop on Natural and Artificial Immune Systems*, pages 10–17.
- 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. (2020). Solving mixed integer programs using neural networks. *CoRR*, abs/2012.13349.
- Nocedal, J. and Wright, S. J. (2006). *Numerical optimization*. Springer, 2 edition.
- Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., and Monfardini, G. (2009). The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80.
- Seyfi, M., Banitalebi-Dehkordi, A., Zhou, Z., and Zhang, Y. (2023). Exact combinatorial optimization with temporo-attentional graph neural networks. In *European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 268–283.
- Shanno, D. F. (2012). Who invented the interior-point method?
- Song, J., Lanka, R., Yue, Y., and Dilkina, B. (2020). A general large neighborhood search framework for solving integer linear programs. In *Advances in Neural Information Processing Systems*.
- Sperduti, A. and Starita, A. (1997). Supervised neural networks for the classification of structures. *IEEE Transactions on Neural Networks*, 8(3):714–35.
- Veličković, P., Ying, R., Padovano, M., Hadsell, R., and Blundell, C. (2020). Neural execution of graph algorithms. In *International Conference on Learning Representations*.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. (2018). Graph attention networks. In *International Conference on Learning Representations*.
- Wu, D. and Lisser, A. (2023). A deep learning approach for solving linear programming problems. *Neurocomputing*, 520:15–24.
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. (2019). How powerful are graph neural networks? In *International Conference on Learning Representations*.
- Zarpellon, G., Jo, J., Lodi, A., and Bengio, Y. (2020). Parameterizing branch-and-bound search trees to learn branching policies. *CoRR*, abs/2002.05120.

Checklist

- For all models and algorithms presented, check if you include:
 - A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
 - (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes]
- For any theoretical claim, check if you include:
 - Statements of the full set of assumptions of all theoretical results. [Yes]
 - Complete proofs of all theoretical results. [Yes]
 - Clear explanations of any assumptions. [Yes]
- For all figures and tables that present empirical results, check if you include:

- (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
- (a) Citations of the creator If your work uses existing assets. [Not Applicable]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Yes]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
- (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Algorithm 3 Conjugate gradient algorithm for IPMs

Input: An instance $I = (\mathbf{A}, \mathbf{b}, \mathbf{c})$ and a point $(\mathbf{x}, \mathbf{w}, \mathbf{s}, \mathbf{r}) > \mathbf{0}$.

- 1: $\mathbf{p} \leftarrow \mathbf{b} - \mathbf{A}\mathbf{x} + \sigma\mu\mathbf{D}(\mathbf{w})^{-1}\mathbf{1} + \mathbf{A}\mathbf{D}(\mathbf{s})^{-1}\mathbf{D}(\mathbf{x})[\mathbf{c} - \mathbf{A}^\top\mathbf{w} - \sigma\mu\mathbf{D}(\mathbf{x})^{-1}\mathbf{1}]$
- 2: $\mathbf{v} \leftarrow -\mathbf{p}$
- 3: $\Delta\mathbf{w} \leftarrow \mathbf{0}$
- 4: **for** m iterations **do**
- 5: $\mathbf{u} \leftarrow [\mathbf{A}\mathbf{D}(\mathbf{s})^{-1}\mathbf{D}(\mathbf{x})\mathbf{A}^\top + \mathbf{D}(\mathbf{w})^{-1}\mathbf{D}(\mathbf{r})]\mathbf{p}$
- 6: $\alpha \leftarrow \mathbf{v}^\top\mathbf{v}/\mathbf{p}^\top\mathbf{u}$
- 7: $\Delta\mathbf{w} += \alpha\mathbf{p}$
- 8: $\mathbf{v}_{\text{new}} \leftarrow \mathbf{v} + \alpha\mathbf{u}$
- 9: $\beta \leftarrow \mathbf{v}_{\text{new}}^\top\mathbf{v}_{\text{new}}/\mathbf{v}^\top\mathbf{v}$
- 10: $\mathbf{v} \leftarrow \mathbf{v}_{\text{new}}$
- 11: $\mathbf{p} \leftarrow -\mathbf{v} + \beta\mathbf{p}$
- 12: **end for**
- 13: **return** A direction $\Delta\mathbf{w}$ that solves the system (6).

A Missing proofs

In this section, we prove Theorem 1 and Proposition 2. To do so, first notice the following. In Step 3 of Algorithm 1 and Step 2 of Algorithm 2, we must solve the system Equation (6). This can be done as follows. Since $(\mathbf{x}, \mathbf{w}, \mathbf{s}, \mathbf{r}) \geq \mathbf{0}$ at any step of the algorithm, the matrix $\mathbf{Q} = \mathbf{A}\mathbf{D}(\mathbf{s})^{-1}\mathbf{D}(\mathbf{x})\mathbf{A}^\top + \mathbf{D}(\mathbf{w})^{-1}\mathbf{D}(\mathbf{r})$ is always symmetric and positive definite. Therefore, we can solve the system with the conjugate gradient algorithm (Nocedal and Wright, 2006, Algorithm 5.2), say with initial point set at $\Delta\mathbf{w}_0 = \mathbf{0}$ for simplicity. A specialization of the algorithm to the problem of solving Equation (6) is described as Algorithm 3.

Lemma 3. There exists a MPNN $f_{\text{MPNN,CG}}$ composed of a $\mathcal{O}(m)$ successive message-passing steps that reproduces Algorithm 3, in the sense which for any LP instance $I = (\mathbf{A}, \mathbf{b}, \mathbf{c})$ and any point $(\mathbf{x}, \mathbf{s}, \mathbf{w}, \mathbf{r}) > \mathbf{0}$, $f_{\text{MPNN,CG}}$ maps the graph $G(I)$ carrying $[\mathbf{x}, \mathbf{s}]$ on the variable nodes, $[\mathbf{w}, \mathbf{r}]$ on the constraint nodes and $[\mu]$ on the objective node to the same graph $G(I)$ carrying the output $[\Delta\mathbf{w}]$ of Algorithm 3 on the constraint nodes.

Proof. We will go through every step of the algorithm and show that it can be computed by message-passing steps on $G(I)$.

- For Step 1, the computation can be broken down as follows. First, we can compute $\mathbf{h}_1 \leftarrow \mathbf{A}^\top\mathbf{w}$, $\mathbf{h}_2 \leftarrow \mu\mathbf{1}_n$ and $\mathbf{h}_3 \leftarrow \mathbf{c}$ by a constraints-to-variables and two objective-to-variables message-passing steps, respectively. Next, one can compute $\mathbf{h}_4 = -\mathbf{x} + \mathbf{D}(\mathbf{s})^{-1}\mathbf{D}(\mathbf{x})[\mathbf{h}_3 - \mathbf{h}_1 - \sigma\mathbf{D}(\mathbf{x})^{-1}\mathbf{h}_2]$ as a local operation on variable nodes. Then we can compute $\mathbf{h}_5 \leftarrow \mathbf{A}\mathbf{h}_4$, $\mathbf{h}_6 \leftarrow \mu\mathbf{1}_m$ and $\mathbf{h}_7 \leftarrow \mathbf{b}$ as a variables-to-constraints and two objective-to-constraints message-passing steps, respectively. Finally, we can compute $\mathbf{p} \leftarrow \mathbf{h}_7 + \mathbf{h}_5 + \sigma\mathbf{D}(\mathbf{w})^{-1}\mathbf{h}_6$ as a local operation on constraint nodes.
- Steps 2 and 3 are just local operations on constraint nodes.
- Step 5 can be broken down as $\mathbf{h}_1 \leftarrow \mathbf{A}^\top\mathbf{p}$, $\mathbf{h}_2 \leftarrow \mathbf{D}(\mathbf{s})^{-1}\mathbf{D}(\mathbf{x})\mathbf{h}_1$, $\mathbf{h}_3 \leftarrow \mathbf{A}\mathbf{h}_2$, $\mathbf{u} \leftarrow \mathbf{h}_3 + \mathbf{D}(\mathbf{w})^{-1}\mathbf{D}(\mathbf{r})\mathbf{p}$. This can be realized as a constraints-to-variables message-passing step, a local operation on variable nodes, a variables-to-constraints message-passing step, and a local operation on constraint nodes.
- Step 6 can be broken down as $\mathbf{h}_1 \leftarrow \mathbf{v}^\top\mathbf{v}$, $\mathbf{h}_2 \leftarrow \mathbf{p}^\top\mathbf{u}$, $\alpha \leftarrow \mathbf{h}_1/\mathbf{h}_2$. This can be realized as a constraints-to-objective message-passing step, another constraints-to-objective message-passing step, and a local operation on the objective node.
- Step 7 can be broken down as a message-passing step from the objective node to the constraint nodes $\mathbf{h}_1 \leftarrow \alpha\mathbf{1}$, followed by a local operation on the constraint nodes $\Delta\mathbf{w} \leftarrow \Delta\mathbf{w} + \mathbf{D}(\mathbf{h}_1)\mathbf{p}$.
- Similarly, step 8 can be written an objective-to-constraints message-passing step $\mathbf{h}_1 \leftarrow \alpha\mathbf{1}$, followed by a local operation on constraint nodes $\Delta\mathbf{v}_{\text{new}} \leftarrow \mathbf{v} + \mathbf{D}(\mathbf{h}_1)\mathbf{u}$.
- Step 9 can be broken down as $\mathbf{h}_1 \leftarrow \mathbf{v}_{\text{new}}^\top\mathbf{v}_{\text{new}}$, $\mathbf{h}_2 \leftarrow \mathbf{v}^\top\mathbf{v}$, $\beta \leftarrow \mathbf{h}_1/\mathbf{h}_2$. This can be realized as a constraints-to-objective message-passing step, another constraints-to-objective message-passing step, and a local operation on the objective node.

- Step 10 is a local operation on constraint nodes.
- Finally, step 11 can be written an objective-to-constraints message-passing step $\mathbf{h}_1 \leftarrow \beta \mathbf{1}$, followed by a local operation on constraint nodes $\Delta \mathbf{p} \leftarrow -\mathbf{v} + \mathbf{D}(\mathbf{h}_1)\mathbf{p}$.

Counting the number of successive message-passing steps, we find that Steps 1–3 can be realized in 8 message-passing steps, while each iteration, comprised of Steps 5–11, can be realized in 9 message-passing steps, completing the proof. \square

We now move on with the proofs of Theorem 1 and Proposition 2.

Theorem 4. There exists an MPNN $f_{\text{MPNN,IPM1}}$ composed of $\mathcal{O}(m)$ message-passing steps that reproduces an iteration of Algorithm 1, in the sense that for any LP instance $I = (\mathbf{A}, \mathbf{b}, \mathbf{c})$ and any iteration step $t \geq 0$, $f_{\text{MPNN,IPM1}}$ maps the graph $G(I)$ carrying $[\mathbf{x}_t, \mathbf{s}_t]$ on the variable nodes and $[\mathbf{w}_t, \mathbf{r}_t]$ on the constraint nodes to the same graph $G(I)$ carrying $[\mathbf{x}_{t+1}, \mathbf{s}_{t+1}]$ on the variable nodes and $[\mathbf{w}_{t+1}, \mathbf{r}_{t+1}]$ on the constraint nodes.

Proof. We need to check that every step can be computed by message-passing steps over $G(I)$.

- Step 2 can be written as $\mathbf{h}_1 \leftarrow \mathbf{x}^\top \mathbf{s}$, $\mathbf{h}_2 \leftarrow \mathbf{w}^\top \mathbf{r}$, $\mu = (\mathbf{h}_1 + \mathbf{h}_2)/(n + m)$. These can be realized as a variable-to-objective message-passing step, a constraints-to-objective message-passing step, and a local operation on the objective node, respectively.
- Step 3 can be written as message-passing steps by Lemma 3.
- Step 4 can be broken down as follows. We can compute $\mathbf{h}_1 \leftarrow \mathbf{A}^\top[\mathbf{w} + \Delta \mathbf{w}]$, $\mathbf{h}_2 \leftarrow \mu \mathbf{1}_n$ and $\mathbf{h}_3 \leftarrow \mathbf{c}$ by a constraints-to-variables and two objective-to-variables message-passing steps, respectively. Then, one can compute $\Delta \mathbf{x} \leftarrow \mathbf{D}(\mathbf{s})^{-1} \mathbf{D}(\mathbf{x})\mathbf{h}_1 - \mathbf{h}_3 + \sigma \mathbf{D}(\mathbf{x})^{-1} \mathbf{h}_2$ by a local operation on variable nodes.
- Step 5 can be realized by taking an objective-to-variables message-passing step $\mathbf{h}_1 \leftarrow \mu \mathbf{1}_n$, and computing $\Delta \mathbf{s} \leftarrow \sigma \mathbf{D}(\mathbf{x})^{-1} \mathbf{h}_1 - \mathbf{s} - \mathbf{D}(\mathbf{x})^{-1} \mathbf{D}(\mathbf{s})\Delta \mathbf{x}$.
- Step 6 can be realized by taking an objective-to-constraints message-passing step $\mathbf{h}_1 \leftarrow \mu \mathbf{1}_m$, and computing $\Delta \mathbf{r} \leftarrow \sigma \mathbf{D}(\mathbf{w})^{-1} \mathbf{h}_1 - \mathbf{r} - \mathbf{D}(\mathbf{w})^{-1} \mathbf{D}(\mathbf{r})\Delta \mathbf{w}$.
- Step 7 can be performed by message-passing steps as follows. We need to find the largest $\alpha \in (0, 1)$ such that

$$(\mathbf{x} + \alpha \Delta \mathbf{x})_i (\mathbf{s} + \alpha \Delta \mathbf{s})_i \tag{8}$$

$$\geq \gamma \frac{(\mathbf{x} + \alpha \Delta \mathbf{x})^\top (\mathbf{s} + \alpha \Delta \mathbf{s}) + (\mathbf{w} + \alpha \Delta \mathbf{w})^\top (\mathbf{r} + \alpha \Delta \mathbf{r})}{n + m} \tag{9}$$

for every $i \in [n]$ and

$$(\mathbf{w} + \alpha \Delta \mathbf{w})_j (\mathbf{r} + \alpha \Delta \mathbf{r})_j \tag{10}$$

$$\geq \gamma \frac{(\mathbf{x} + \alpha \Delta \mathbf{x})^\top (\mathbf{s} + \alpha \Delta \mathbf{s}) + (\mathbf{w} + \alpha \Delta \mathbf{w})^\top (\mathbf{r} + \alpha \Delta \mathbf{r})}{n + m} \tag{11}$$

for every $j \in [m]$. Equivalently, for each $i \in [n]$, we can find the largest $\alpha_i < 1$ such that 9 holds, that is such that

$$\begin{aligned} & \alpha_i^2 \left(\Delta x_i \Delta s_i - \gamma \frac{\Delta \mathbf{x}^\top \Delta \mathbf{s} + \Delta \mathbf{w}^\top \Delta \mathbf{r}}{n + m} \right) \\ & + \alpha_i \left(x_i \Delta s_i + \Delta x_i s_i - \gamma \frac{\mathbf{x}^\top \Delta \mathbf{s} + \Delta \mathbf{x}^\top \mathbf{s} + \mathbf{w}^\top \Delta \mathbf{r} + \Delta \mathbf{w}^\top \mathbf{r}}{n + m} \right) \\ & + \left(x_i s_i - \gamma \frac{\mathbf{x}^\top \mathbf{s} + \mathbf{w}^\top \mathbf{r}}{n + m} \right) \geq 0 \end{aligned}$$

holds; and similarly, find the largest $\bar{\alpha}_j < 1$ such that 11 holds, that is such that

$$\begin{aligned} & \bar{\alpha}_j^2 \left(\Delta w_j \Delta r_j - \gamma \frac{\Delta \mathbf{x}^\top \Delta \mathbf{s} + \Delta \mathbf{w}^\top \Delta \mathbf{r}}{n+m} \right) \\ & + \bar{\alpha}_j \left(w_j \Delta r_j + \Delta w_j r_j - \gamma \frac{\mathbf{x}^\top \Delta \mathbf{s} + \Delta \mathbf{x}^\top \mathbf{s} + \mathbf{w}^\top \Delta \mathbf{r} + \Delta \mathbf{w}^\top \mathbf{r}}{n+m} \right) \\ & + \left(w_j r_j - \gamma \frac{\mathbf{x}^\top \mathbf{s} + \mathbf{w}^\top \mathbf{r}}{n+m} \right) \geq 0 \end{aligned}$$

holds; then $\alpha = \min_{i,j} \{\alpha_i, \bar{\alpha}_j\}$.

This can be computed by message-passing steps as follows. First, we can compute $\mathbf{h}_1 \leftarrow \Delta \mathbf{x}^\top \Delta \mathbf{s}$, $\mathbf{h}_2 \leftarrow \Delta \mathbf{x}^\top \mathbf{s}$, $\mathbf{h}_3 \leftarrow \mathbf{x}^\top \Delta \mathbf{s}$ and $\mathbf{h}_4 \leftarrow \mathbf{x}^\top \mathbf{s}$ by variable-to-objective message-passing steps; and similarly $\bar{\mathbf{h}}_1 \leftarrow \Delta \mathbf{w}^\top \Delta \mathbf{r}$, $\bar{\mathbf{h}}_2 \leftarrow \Delta \mathbf{w}^\top \mathbf{r}$, $\bar{\mathbf{h}}_3 \leftarrow \mathbf{w}^\top \Delta \mathbf{r}$, $\bar{\mathbf{h}}_4 \leftarrow \mathbf{w}^\top \mathbf{r}$ by constraints-to-objective message-passing steps. The quantities $t_1 \leftarrow \gamma(\mathbf{h}_1 + \bar{\mathbf{h}}_1)/(n+m)$, $t_2 \leftarrow \gamma(\mathbf{h}_2 + \mathbf{h}_3 + \bar{\mathbf{h}}_2 + \bar{\mathbf{h}}_3)/(n+m)$ and $t_3 \leftarrow \gamma(\mathbf{h}_4 + \bar{\mathbf{h}}_4)/(n+m)$ can then be computed by local operations on the objective node, and returned to the variable and constraint nodes by objective-to-variables message-passing steps $\mathbf{t}_1 \leftarrow t_1 \mathbf{1}_n$, $\mathbf{t}_2 \leftarrow t_2 \mathbf{1}_n$, $\mathbf{t}_3 \leftarrow t_3 \mathbf{1}_n$ and objective-to-constraint message-passing steps $\bar{\mathbf{t}}_1 \leftarrow t_1 \mathbf{1}_n$, $\bar{\mathbf{t}}_2 \leftarrow t_2 \mathbf{1}_n$, $\bar{\mathbf{t}}_3 \leftarrow t_3 \mathbf{1}_n$. Then, on each variable node v_i , we can solve

$$\begin{aligned} \alpha_i = \max \{ & \alpha \in (0, 1) \mid \alpha^2 (\Delta x_i \Delta s_i - t_{1i}) \\ & + \alpha (x_i \Delta s_i + \Delta x_i s_i - t_{2i}) + (x_i s_i - t_{3i}) \geq 0 \} \end{aligned}$$

as a local operation, and similarly, on each constraint node c_j , we can find

$$\begin{aligned} \bar{\alpha}_j = \max \{ & \alpha \in (0, 1) \mid \alpha^2 (\Delta w_j \Delta r_j - \bar{t}_{1j}) \\ & + \alpha (w_j \Delta r_j + \Delta w_j r_j - \bar{t}_{2j}) + (w_j r_j - \bar{t}_{3j}) \geq 0 \} \end{aligned}$$

as a local operation. Finally, we can compute $\alpha_v \leftarrow \min_i \alpha_i$ as a variables-to-objective message-passing step, and $\alpha_c \leftarrow \min_j \bar{\alpha}_j$ as a constraints-to-objective message-passing step, and finally take $\alpha = \min(\alpha_v, \alpha_c)$ as a local operation on the objective node.

- Finally, step 8 can be performed by taking objective-to-variables, objective-to-constraints message-passing steps $\mathbf{h}_1 \leftarrow \alpha \mathbf{1}_n$, $\bar{\mathbf{h}}_1 \leftarrow \alpha \mathbf{1}_m$, and taking local operations $(\mathbf{x}, \mathbf{s}) \leftarrow (\mathbf{x} + \mathbf{D}(\mathbf{h}_1) \Delta \mathbf{x}, \mathbf{s} + \mathbf{D}(\mathbf{h}_1) \Delta \mathbf{s})$ and $(\mathbf{w}, \mathbf{r}) \leftarrow (\mathbf{w} + \mathbf{D}(\bar{\mathbf{h}}_1) \Delta \mathbf{w}, \mathbf{r} + \mathbf{D}(\bar{\mathbf{h}}_1) \Delta \mathbf{r})$ on variable and constraint nodes respectively.

Counting the number of successive message-passing steps, we find that all steps can be realized in 23 message-passing steps, plus the $\mathcal{O}(m)$ steps of Step 3, completing the proof. \square

Further, we show an analogous result for Algorithm 2.

Proposition 5. There exists an MPNN $f_{\text{MPNN,IPM2}}$ composed of $\mathcal{O}(m)$ message-passing steps that reproduces each iteration of Algorithm 2, in the sense that for any LP instance $I = (\mathbf{A}, \mathbf{b}, \mathbf{c})$ and any iteration step $t \geq 0$, $f_{\text{MPNN,IPM2}}$ maps the graph $G(I)$ carrying $[\mathbf{x}_t, \mathbf{s}_t]$ on the variable nodes, $[\mathbf{w}_t, \mathbf{r}_t]$ on the constraint nodes and $[\mu_t]$ on the objective node to the same graph $G(I)$ carrying $[\mathbf{x}_{t+1}, \mathbf{s}_{t+1}]$ on the variable nodes, $[\mathbf{w}_{t+1}, \mathbf{r}_{t+1}]$ on the constraint nodes and $[\mu_{t+1}]$ on the objective node.

Proof. We need to show that every step can be written as message-passing steps over $G(I)$, as in Theorem 1. Steps 2-5 are the same as 3-6, so by Theorem 1, they can be written as message-passing steps. This leaves to check steps 6-8.

- The analysis of Step 6 is similar to the analysis of Step 7 of Algorithm 1 in Theorem 1, but simpler. On every variable node, we can compute

$$\begin{aligned} \alpha_i = \max \{ & \alpha \in (0, \infty) \mid \alpha^2 \Delta x_i \Delta s_i \\ & + \alpha (x_i \Delta s_i + \Delta x_i s_i) + x_i s_i \geq 0 \} \end{aligned}$$

as a local operation, and similarly, on each constraint node c_j we can find

$$\begin{aligned} \bar{\alpha}_j = \max \{ & \alpha \in (0, \infty) \mid \alpha^2 \Delta w_j \Delta r_j \\ & + \alpha (w_j \Delta r_j + \Delta w_j r_j) + w_j r_j \geq 0 \} \end{aligned}$$

as a local operation. Then we can compute $\alpha_v \leftarrow \min_i \alpha_i$ as a variables-to-objective message-passing step, and $\alpha_c \leftarrow \min_j \bar{\alpha}_j$ as a constraints-to-objective message-passing step, and finally take $\alpha = \min(\alpha_v, \alpha_c)$ as a local operation on the objective node.

- The analysis of step 7 is similar to the analysis of Step 8 of Algorithm 1 in Theorem 1 as well. It can be performed by taking objective-to-variables, objective-to-constraints message-passing steps $\mathbf{h}_1 \leftarrow \alpha \mathbf{1}_n$, $\bar{\mathbf{h}}_1 \leftarrow \alpha \mathbf{1}_m$, and taking local operations $(\mathbf{x}, \mathbf{s}) \leftarrow (\mathbf{x} + 0.99\mathbf{D}(\mathbf{h}_1)\Delta\mathbf{x}, \mathbf{s} + 0.99\mathbf{D}(\mathbf{h}_1)\Delta\mathbf{s})$ and $(\mathbf{w}, \mathbf{r}) \leftarrow (\mathbf{w} + 0.99\mathbf{D}(\bar{\mathbf{h}}_1)\Delta\mathbf{w}, \mathbf{r} + 0.99\mathbf{D}(\bar{\mathbf{h}}_1)\Delta\mathbf{r})$ on variable and constraint nodes respectively.
- Finally, step 8 is a local operation on the objective node.

Just like in Theorem 4, we can see all the computations can be realized in $\mathcal{O}(m)$ message-passing steps, completing the proof. \square

B Details of IPM-MPNNs

In the following, we outline details regarding the specific MPNN layers used in Section 5. We follow the notation outlined in Section 3. Furthermore, let MLP be a multi-layer perceptron whose subscript denotes its role. Specifically, MLP_* is for node initialization or node updating after gathering message functions, $\text{MLP}_{* \rightarrow *}$ is for message vector mapping, and MLP_{**} is for edge feature embedding in each layer. At the initialization $t = 0$, we obtain node embeddings by

$$\begin{aligned}\mathbf{h}_v^{(0)} &:= \text{MLP}_v^{(0)}(\mathbf{x}_v), \forall v \in V(I), \\ \mathbf{h}_c^{(0)} &:= \text{MLP}_c^{(0)}(\mathbf{x}_c), \forall c \in C(I), \\ \mathbf{h}_o^{(0)} &:= \text{MLP}_o^{(0)}(\mathbf{x}_o).\end{aligned}$$

Then, a GCN layer updates the constraint, objective, and variable nodes as follows:

$$\begin{aligned}\mathbf{h}_c^{(t)} &:= \text{MLP}_c^{(t)} \left[\text{MLP}_{c \rightarrow c}^{(t)}(\mathbf{h}_c^{(t-1)}) + \right. \\ &\quad \left. \text{MLP}_{o \rightarrow c}^{(t)} \left(\frac{1}{\sqrt{d_o d_c}} (\mathbf{h}_o^{(t-1)} + \text{MLP}_{oc}^{(t)}(\mathbf{e}_{oc})) \right) + \right. \\ &\quad \left. \text{MLP}_{v \rightarrow c}^{(t)} \left(\sum_{v \in N_c \cap V(I)} \frac{1}{\sqrt{d_v d_c}} (\mathbf{h}_v^{(t-1)} + \text{MLP}_{vc}^{(t)}(\mathbf{e}_{vc})) \right) \right] \\ \mathbf{h}_o^{(t)} &:= \text{MLP}_o^{(t)} \left[\text{MLP}_{o \rightarrow o}^{(t)}(\mathbf{h}_o^{(t-1)}) + \right. \\ &\quad \left. \text{MLP}_{c \rightarrow o}^{(t)} \left(\sum_{c \in C(I)} \frac{1}{\sqrt{d_o d_c}} (\mathbf{h}_c^{(t-1)} + \text{MLP}_{co}^{(t)}(\mathbf{e}_{co})) \right) + \right. \\ &\quad \left. \text{MLP}_{v \rightarrow o}^{(t)} \left(\sum_{v \in V(I)} \frac{1}{\sqrt{d_o d_v}} (\mathbf{h}_v^{(t-1)} + \text{MLP}_{vo}^{(t)}(\mathbf{e}_{vo})) \right) \right] \\ \mathbf{h}_v^{(t)} &:= \text{MLP}_v^{(t)} \left[\text{MLP}_{v \rightarrow v}^{(t)}(\mathbf{h}_v^{(t-1)}) + \right. \\ &\quad \left. \text{MLP}_{o \rightarrow v}^{(t)} \left(\frac{1}{\sqrt{d_o d_c}} (\mathbf{h}_o^{(t-1)} + \text{MLP}_{ov}^{(t)}(\mathbf{e}_{ov})) \right) + \right. \\ &\quad \left. \text{MLP}_{c \rightarrow v}^{(t)} \left(\sum_{c \in N_v \cap C(I)} \frac{1}{\sqrt{d_c d_v}} (\mathbf{h}_c^{(t-1)} + \text{MLP}_{cv}^{(t)}(\mathbf{e}_{cv})) \right) \right].\end{aligned}$$

Similarly, for the GIN layer, we have:

$$\begin{aligned}
 \mathbf{h}_c^{(t)} &:= \text{MLP}_c^{(t)} \left[\left((1 + \epsilon_c^{(t)}) \text{MLP}_{c \rightarrow c}^{(t)}(\mathbf{h}_c^{(t-1)}) + \right. \right. \\
 &\quad \left. \text{MLP}_{o \rightarrow c}^{(t)}(\mathbf{h}_o^{(t-1)} + \text{MLP}_{oc}^{(t)}(\mathbf{e}_{oc})) + \right. \\
 &\quad \left. \left. \text{MLP}_{v \rightarrow c}^{(t)} \left(\sum_{v \in N_c \cap V(I)} (\mathbf{h}_v^{(t-1)} + \text{MLP}_{vc}^{(t)}(\mathbf{e}_{vc})) \right) \right) \right] \\
 \mathbf{h}_o^{(t)} &:= \text{MLP}_o^{(t)} \left[\left((1 + \epsilon_o^{(t)}) \text{MLP}_{o \rightarrow o}^{(t)}(\mathbf{h}_o^{(t-1)}) + \right. \right. \\
 &\quad \left. \text{MLP}_{c \rightarrow o}^{(t)} \left(\sum_{c \in C(I)} (\mathbf{h}_c^{(t)} + \text{MLP}_{co}^{(t)}(\mathbf{e}_{co})) \right) + \right. \\
 &\quad \left. \left. \text{MLP}_{v \rightarrow o}^{(t)} \left(\sum_{v \in V(I)} (\mathbf{h}_v^{(t-1)} + \text{MLP}_{vo}^{(t)}(\mathbf{e}_{vo})) \right) \right) \right] \\
 \mathbf{h}_v^{(t)} &:= \text{MLP}_v^{(t)} \left[\left((1 + \epsilon_v^{(t)}) \text{MLP}_{v \rightarrow v}^{(t)}(\mathbf{h}_v^{(t-1)}) + \right. \right. \\
 &\quad \left. \text{MLP}_{o \rightarrow v}^{(t)}(\mathbf{h}_o^{(t)} + \text{MLP}_{ov}^{(t)}(\mathbf{e}_{ov})) + \right. \\
 &\quad \left. \left. \text{MLP}_{c \rightarrow v}^{(t)} \left(\sum_{c \in N_v \cap C(I)} (\mathbf{h}_c^{(t)} + \text{MLP}_{cv}^{(t)}(\mathbf{e}_{cv})) \right) \right) \right].
 \end{aligned}$$

For the GEN layer, we have:

$$\begin{aligned}
 \mathbf{h}_c^{(t)} &:= \text{MLP}_c^{(t)} \left[\text{MLP}_{c \rightarrow c}^{(t)}(\mathbf{h}_c^{(t-1)}) + \right. \\
 &\quad \left. \text{MLP}_{o \rightarrow c}^{(t)}(\mathbf{h}_o^{(t-1)} + \text{MLP}_{oc}^{(t)}(\mathbf{e}_{oc}) + \epsilon_{o \rightarrow c}^{(t)}) + \right. \\
 &\quad \left. \text{MLP}_{v \rightarrow c}^{(t)} \left(\text{MSG} \left(\left\{ \left\{ \mathbf{h}_v^{(t-1)} + \text{MLP}_{vc}^{(t)}(\mathbf{e}_{vc}) + \epsilon_{v \rightarrow c}^{(t)} \mid v \in N_c \cap V(I) \right\} \right\} \right) \right) \right] \\
 \mathbf{h}_o^{(t)} &:= \text{MLP}_o^{(t)} \left[\text{MLP}_{o \rightarrow o}^{(t)}(\mathbf{h}_o^{(t-1)}) + \right. \\
 &\quad \left. \text{MLP}_{c \rightarrow o}^{(t)} \left(\text{MSG} \left(\left\{ \left\{ \mathbf{h}_c^{(t-1)} + \text{MLP}_{co}^{(t)}(\mathbf{e}_{co}) + \epsilon_{c \rightarrow o}^{(t)} \mid c \in C(I) \right\} \right\} \right) \right) + \right. \\
 &\quad \left. \text{MLP}_{v \rightarrow o}^{(t)} \left(\text{MSG} \left(\left\{ \left\{ \mathbf{h}_v^{(t-1)} + \text{MLP}_{vo}^{(t)}(\mathbf{e}_{vo}) + \epsilon_{v \rightarrow o}^{(t)} \mid v \in V(I) \right\} \right\} \right) \right) \right] \\
 \mathbf{h}_v^{(t)} &:= \text{MLP}_v^{(t)} \left[\text{MLP}_{v \rightarrow v}^{(t)}(\mathbf{h}_v^{(t-1)}) + \right. \\
 &\quad \left. \text{MLP}_{o \rightarrow v}^{(t)}(\mathbf{h}_o^{(t)} + \text{MLP}_{ov}^{(t)}(\mathbf{e}_{ov}) + \epsilon_{o \rightarrow v}^{(t)}) + \right. \\
 &\quad \left. \text{MLP}_{c \rightarrow v}^{(t)} \left(\text{MSG} \left(\left\{ \left\{ \mathbf{h}_c^{(t-1)} + \text{MLP}_{cv}^{(t)}(\mathbf{e}_{cv}) + \epsilon_{c \rightarrow v}^{(t)} \mid c \in N_v \cap C(I) \right\} \right\} \right) \right) \right],
 \end{aligned}$$

where MSG is the softmax aggregation with $\tau = 1$, i.e.,

$$\text{softmax}(\mathcal{X} \mid \tau) = \sum_{\mathbf{x}_i \in \mathcal{X}} \frac{\exp(\tau \cdot \mathbf{x}_i)}{\sum_{\mathbf{x}_j \in \mathcal{X}} \exp(\tau \cdot \mathbf{x}_j)} \cdot \mathbf{x}_i.$$

C Details of datasets

In the following, we describe our datasets.

C.1 Combinatorial optimization problems

In the following, we briefly describe the combinatorial optimization problems.

Set cover problem The set cover problem aims at covering the universe U with a collection of given subsets $S = \{S_1, S_2, \dots, S_m\}$ satisfying $\cup_{i=1}^m S_i = U$, with the target of minimizing the objective function. Formally, let x_i be the variable deciding whether subset S_i is selected, and c_i the cost per subset, we have:

$$\begin{aligned} \min_{\mathbf{x}} \quad & \sum_{i=1}^m c_i x_i \\ \text{s.t.} \quad & \sum_{i:u \in S_i} x_i \geq 1, \forall u \in U \\ & x_i \in \{0, 1\}, \forall i \in [m]. \end{aligned}$$

Maximal independent set Given an undirected graph G with node set $V(G)$ and edge set $E(G)$, the goal of the maximal independent set problem is to find a set of nodes where no pairs of them are connected. If we use x_i to denote a node i is selected or not, we have:

$$\begin{aligned} \max_{\mathbf{x}} \quad & \sum_{i \in V} x_i \\ \text{s.t.} \quad & x_u + x_v \leq 1, \forall (u, v) \in E(G), u, v \in V(G) \\ & x_i \in \{0, 1\}, \forall i \in V(G). \end{aligned}$$

Combinatorial auction Suppose there are a set of items M and bidders N . Each bidder $i \in N$ maintains a set of bids B_i , each bid $b \in B_i$ is associated with a subset $S_{ib} \subseteq M$ of items and a value v_{ib} that the bidder i is willing to pay for this subset. The binary decision variable x_{ib} is 1 if the bid b by bidder i is accepted or 0 otherwise. The MILP formulation of the problem is as follows:

$$\begin{aligned} \max_{\mathbf{x}} \quad & \sum_{i \in N} \sum_{b \in B_i} v_{ib} x_{ib} \\ \text{s.t.} \quad & \sum_{i \in N} \sum_{b \in B_i: j \in S_{ib}} x_{ib} \leq 1, \forall j \in M, \\ & x_{ib} \in \{0, 1\}, \forall i \in N, b \in B_i. \end{aligned} \tag{12}$$

Capacitated facility location Given a set of facilities M and another of customers N , we aim to build facilities satisfying the customers' demand at minimum cost. Let $y_j, j \in M$ be the binary decision of building the facility j or not, and x_{ij} be a continuous variable indicating the fraction of demand facility j sends to customer $i \in N$. Let $d_i \in \mathbb{R}^+$ be the amount of demand of customer i , and v_j be the volume of the facility j , c_{ij} the cost of shipment, and f_j the cost of building facility j , we formulate the problem as follows:

$$\begin{aligned} \min_{\mathbf{x}} \quad & \sum_{j \in M} f_j y_j + \sum_{i \in N} \sum_{j \in M} c_{ij} x_{ij} \\ \text{s.t.} \quad & \sum_{j \in M} x_{ij} = 1, \forall i \in N \\ & \sum_{i \in N} d_i x_{ij} \leq v_j y_j, \forall i \in N, j \in M \\ & y_j \in \{0, 1\}, x_{ij} \in [0, 1], \forall i \in N, j \in M. \end{aligned} \tag{13}$$

Table 5: Sizes of setcover instances.

| Size | Num. Row | Num. Col. | Density | Num. instances |
|-------|------------|------------|---------|----------------|
| Mini | [15, 20] | [15, 20] | 0.15 | 1000 |
| Small | [30, 50] | [50, 70] | 0.05 | 10 000 |
| Large | [300, 500] | [500, 700] | 0.01 | 10 000 |

Table 6: Sizes of maximal independent set instances.

| Size | Num. nodes | Affinity | Num. instances |
|-------|------------|----------|----------------|
| Mini | [10, 20] | 2 | 1000 |
| Small | [50, 70] | 2 | 10 000 |
| Large | [500, 700] | 2 | 10 000 |

C.2 Generation of instances

We propose various sizes of generated instances; see Tables 5 to 8 for the size parameters of each dataset. The generation of instances follows the setting of Gasse et al. (2019). For the set covering instances, problems are generated with 15-20 rows and columns and a constraint matrix density of 0.15 for mini instances. For small instances, we used 30-50 rows and 50-70 columns with a density of 0.05. Large instances have 300-500 rows and 500-700 columns with a density of 0.01. We employ the Erdős–Rényi random graph as the foundational graph when generating maximal independent set instances, designating 10-20 nodes for mini instances, 50-70 nodes for small instances, and 500-700 nodes for large instances. For combinatorial auction instances, we modulate the size by varying the number of items and bids: specifically, 20 items and bids are set for mini instances, 50-80 for small instances, and 500-800 for large instances. Lastly, for the capacitated facility location instances, we set 3-5 customers and facilities for mini instances, 10 for both in small instances, and 20-30 in large ones.

D Training parameters

For all the experiments, we train the neural networks with Adam optimizer with default hyperparameters, and run for at most 1000 epochs. During training, we leverage learning rate decay with right to the validation objective gap metric with a decay ratio of 0.5 and patience 50. We terminate the run at patience 100 epochs. Besides, we display the other task-specific hyperparameters in Table 9, which are the batch size, number of MPNN layers as well as the number of sampled IPM solver steps, the step decay factor introduced in Equation (7), the loss weight combination in Section 4, plus the weight decay of the optimizer.

With regard to the bipartiteness ablation study, we also tune the hyperparameters for the sake of fair comparison. The chosen hyperparameters are listed in Table 10. Moreover, the hyperparameter configurations for the ODE approach baseline (Wu and Lissner, 2023) are shown in Table 11.

E Extended experimental results

We provide the training time per epoch and maximal GPU memory usage as supplementary results of Table 1 in Table 13. We also show more elaborated solving time in Table 12.

Table 7: Sizes of combinatorial auction instances.

| Size | Num. items | Num. bids | Num. instances |
|-------|------------|------------|----------------|
| Mini | 20 | 20 | 1000 |
| Small | [50, 80] | [50, 80] | 10 000 |
| Large | [300, 500] | [300, 500] | 10 000 |

Table 8: Sizes of capacitated facility location instances.

| Size | Num. customers | Num. facilities | Ratio | Num. instances |
|-------|----------------|-----------------|-------|----------------|
| Mini | [3, 5] | [3, 5] | 5 | 1000 |
| Small | 10 | 10 | 5 | 10 000 |
| Large | [20, 30] | [20, 30] | 5 | 10 000 |

 Table 9: Training hyperparameters of our tripartite MPNN main experiments, v^a represents $v \times 10^a$.

| Instances | Size | MPNN | Batch size | Num. layers | Hidden dim. | α | w_{var} | w_{obj} | w_{cons} | Weight decay |
|-----------|-------|------|------------|-------------|-------------|----------|------------------|------------------|-------------------|--------------|
| Setcover | Small | GEN | 512 | 8 | 180 | 0.2 | 1.2 | 0.8 | 0.2 | 1.2^{-6} |
| | | GCN | 512 | 8 | 180 | 0.8 | 1.0 | 0.3 | 2.2 | 4.4^{-7} |
| | | GIN | 512 | 8 | 180 | 0.7 | 1.0 | 2.4 | 7.5 | 5.6^{-6} |
| | Large | GEN | 128 | 8 | 180 | 0.2 | 1.2 | 0.8 | 0.2 | 1.2^{-6} |
| | | GCN | 128 | 8 | 180 | 0.2 | 1.0 | 2.2 | 0.3 | 1.5^{-8} |
| | | GIN | 128 | 8 | 180 | 0.7 | 1.0 | 4.5 | 2.2 | 2.8^{-8} |
| Indset | Small | GEN | 512 | 8 | 180 | 0.2 | 1.2 | 0.8 | 0.2 | 1.2^{-6} |
| | | GCN | 512 | 8 | 180 | 0.5 | 1.0 | 4.5 | 9.6 | 2.0^{-7} |
| | | GIN | 512 | 8 | 180 | 0.7 | 1.0 | 2.4 | 7.5 | 5.6^{-6} |
| | Large | GEN | 128 | 8 | 180 | 0.2 | 1.2 | 0.8 | 0.2 | 1.2^{-6} |
| | | GCN | 128 | 8 | 180 | 0.5 | 1.0 | 4.5 | 9.6 | 2.0^{-7} |
| | | GIN | 128 | 8 | 180 | 0.7 | 1.0 | 2.4 | 7.5 | 5.6^{-6} |
| Cauc. | Small | GEN | 512 | 8 | 180 | 0.9 | 1.0 | 4.6 | 5.3 | 0.0 |
| | | GCN | 512 | 8 | 180 | 0.4 | 1.0 | 3.4 | 5.8 | 0.0 |
| | | GIN | 512 | 8 | 180 | 0.6 | 1.0 | 4.3 | 6.3 | 0.0 |
| | Large | GEN | 128 | 8 | 180 | 0.9 | 1.0 | 4.6 | 5.3 | 0.0 |
| | | GCN | 128 | 8 | 180 | 0.4 | 1.0 | 3.4 | 5.8 | 0.0 |
| | | GIN | 128 | 8 | 180 | 0.6 | 1.0 | 4.3 | 6.3 | 0.0 |
| Fac. | Small | GEN | 512 | 8 | 180 | 0.8 | 1.0 | 3.0 | 8.2 | 3.8^{-6} |
| | | GCN | 512 | 8 | 96 | 0.6 | 1.0 | 8.7 | 9.6 | 4.5^{-7} |
| | | GIN | 512 | 8 | 180 | 0.8 | 1.0 | 1.3 | 4.6 | 1.0^{-7} |
| | Large | GEN | 128 | 8 | 180 | 0.8 | 1.0 | 3.09 | 8.2 | 3.8^{-6} |
| | | GCN | 128 | 8 | 96 | 0.6 | 1.0 | 8.7 | 9.6 | 4.5^{-7} |
| | | GIN | 128 | 8 | 180 | 0.9 | 1.0 | 2.5 | 4.0 | 1.0^{-5} |

Table 10: Training hyperparameters of our bipartite MPNN ablation experiments, v^a represents $v \times 10^a$.

| Instances | Size | MPNN | Batch size | Num. layers | Hidden dim. | α | w_{var} | w_{obj} | w_{cons} | Weight decay |
|-----------|-------|------|------------|-------------|-------------|----------|------------------|------------------|-------------------|--------------|
| Setcover | Small | GEN | 512 | 8 | 32 | 0.8 | 1.0 | 2.6 | 0.8 | 1.0^{-6} |
| | | GCN | 512 | 8 | 32 | 0.9 | 1.0 | 5.5 | 1.1 | 1.0^{-5} |
| | | GIN | 512 | 8 | 64 | 0.3 | 1.0 | 4.7 | 0.8 | 1.0^{-5} |
| | Large | GEN | 128 | 8 | 32 | 0.8 | 1.0 | 2.6 | 0.8 | 1.0^{-6} |
| | | GCN | 128 | 8 | 32 | 0.9 | 1.0 | 5.5 | 1.1 | 1.0^{-5} |
| | | GIN | 128 | 8 | 64 | 0.3 | 1.0 | 4.7 | 0.8 | 1.0^{-5} |
| Indset | Small | GEN | 512 | 8 | 180 | 0.6 | 1.0 | 4.7 | 2.0 | 0.0 |
| | | GCN | 512 | 8 | 96 | 0.7 | 1.0 | 6.3 | 3.1 | 0.0 |
| | | GIN | 512 | 8 | 180 | 0.7 | 1.0 | 2.4 | 7.5 | 5.6^{-6} |
| | Large | GEN | 128 | 8 | 180 | 0.6 | 1.0 | 4.7 | 2.0 | 0.0 |
| | | GCN | 128 | 8 | 96 | 0.7 | 1.0 | 6.3 | 3.1 | 0.0 |
| | | GIN | 128 | 8 | 180 | 0.7 | 1.0 | 2.4 | 7.5 | 5.6^{-6} |
| Cauc. | Small | GEN | 512 | 8 | 128 | 0.5 | 1.0 | 6.2 | 6.6 | 1.0^{-7} |
| | | GCN | 512 | 8 | 128 | 0.6 | 1.0 | 4.7 | 4.3 | 1.0^{-7} |
| | | GIN | 512 | 8 | 128 | 0.4 | 1.0 | 6.2 | 4.1 | 1.2^{-8} |
| | Large | GEN | 128 | 8 | 128 | 0.5 | 1.0 | 6.2 | 6.6 | 1.0^{-7} |
| | | GCN | 128 | 8 | 128 | 0.6 | 1.0 | 4.7 | 4.3 | 1.0^{-7} |
| | | GIN | 128 | 8 | 128 | 0.4 | 1.0 | 6.2 | 4.1 | 1.2^{-8} |
| Fac. | Small | GEN | 512 | 8 | 128 | 0.9 | 1.0 | 2.9 | 2.5 | 1.0^{-7} |
| | | GCN | 512 | 8 | 96 | 0.7 | 1.0 | 5.3 | 3.8 | 0.0 |
| | | GIN | 512 | 8 | 180 | 0.8 | 1.0 | 1.3 | 4.6 | 1.0^{-7} |
| | Large | GEN | 128 | 8 | 128 | 0.9 | 1.0 | 2.9 | 2.5 | 1.0^{-7} |
| | | GCN | 128 | 8 | 96 | 0.7 | 1.0 | 5.3 | 3.8 | 0.0 |
| | | GIN | 128 | 8 | 180 | 0.8 | 1.0 | 1.3 | 4.6 | 1.0^{-7} |

Table 11: Training hyperparameters of the ODE approach baseline. The experiments are done on the mini-sized instances with GEN-based MPNNs, v^a represents $v \times 10^a$.

| Instances | MPNN | Candidate | Batch size | Num. layers | Hidden dim. | α | w_{var} | w_{obj} | w_{cons} | Weight decay |
|-----------|------|-----------|------------|-------------|-------------|----------|------------------|------------------|-------------------|--------------|
| Setcover | GEN | Ours | 512 | 3 | 128 | 0.3 | 1.0 | 3.5 | 1.3 | 3.3^{-3} |
| | | Baseline | 8 | 3 | 128 | - | - | - | - | 0.0 |
| | GCN | Ours | 512 | 3 | 180 | 0.8 | 1.0 | 6.1 | 1.6 | 1.0^{-6} |
| | | Baseline | 8 | 3 | 180 | - | - | - | - | 0.0 |
| | GIN | Ours | 512 | 3 | 180 | 0.4 | 1.0 | 3.8 | 1.4 | 5.5^{-6} |
| | | Baseline | 8 | 3 | 180 | - | - | - | - | 0.0 |
| Indset | GEN | Ours | 512 | 3 | 128 | 0.4 | 1.0 | 7.1 | 6.2 | 1.0^{-6} |
| | | Baseline | 8 | 3 | 128 | - | - | - | - | 0.0 |
| | GCN | Ours | 512 | 3 | 180 | 0.8 | 1.0 | 3.5 | 5.6 | 2.1^{-6} |
| | | Baseline | 8 | 3 | 180 | - | - | - | - | 0.0 |
| | GIN | Ours | 512 | 3 | 180 | 0.4 | 1.0 | 5.9 | 3.9 | 9.7^{-6} |
| | | Baseline | 8 | 3 | 180 | - | - | - | - | 0.0 |
| Cauc | GEN | Ours | 512 | 3 | 128 | 0.8 | 1.0 | 9.6 | 7.1 | 8.0^{-5} |
| | | Baseline | 8 | 3 | 128 | - | - | - | - | 0.0 |
| | GCN | Ours | 512 | 3 | 180 | 0.7 | 1.0 | 4.5 | 5.0 | 3.7^{-6} |
| | | Baseline | 8 | 3 | 180 | - | - | - | - | 0.0 |
| | GIN | Ours | 512 | 3 | 128 | 0.9 | 1.0 | 6.4 | 5.0 | 1.0^{-6} |
| | | Baseline | 8 | 3 | 128 | - | - | - | - | 0.0 |
| Fac | GEN | Ours | 512 | 3 | 128 | 0.7 | 1.0 | 5.3 | 0.8 | 1.7^{-6} |
| | | Baseline | 8 | 3 | 128 | - | - | - | - | 0.0 |
| | GCN | Ours | 512 | 3 | 128 | 0.8 | 1.0 | 2.9 | 3.7 | 9.2^{-6} |
| | | Baseline | 8 | 3 | 128 | - | - | - | - | 0.0 |
| | GIN | Ours | 512 | 3 | 180 | 0.8 | 1.0 | 1.8 | 1.5 | 3.6^{-7} |
| | | Baseline | 8 | 3 | 180 | - | - | - | - | 0.0 |

Table 12: Comparing IPM-MPNNs' and bipartite GNN ablations' inference time to SciPy's IPM implementation and our Python-based IPM solver. We report mean and standard deviation in seconds over the validation set. We print the best results per target in bold.

| Instances | SciPy Solver | Our Solver | Tripartite | | | Bipartite | | |
|----------------|--------------------|-------------|-------------|-------------|-------------|--------------|--------------|---------------------|
| | | | GEN | GCN | GIN | GEN | GCN | GIN |
| Small setcover | 0.006±0.004 | 0.071±0.015 | 0.033±0.001 | 0.029±0.001 | 0.017±0.001 | 0.011±0.0003 | 0.010±0.0003 | 0.006±0.0003 |
| Large setcover | 0.390±0.098 | 3.696±2.141 | 0.033±0.001 | 0.030±0.001 | 0.021±0.001 | 0.011±0.0002 | 0.011±0.0001 | 0.008±0.0001 |
| Small indset | 0.008±0.007 | 0.089±0.024 | 0.033±0.001 | 0.031±0.002 | 0.021±0.001 | 0.011±0.0001 | 0.011±0.0001 | 0.008±0.0001 |
| Large indset | 0.226±0.087 | 1.053±0.281 | 0.033±0.002 | 0.030±0.001 | 0.021±0.001 | 0.011±0.0001 | 0.011±0.0001 | 0.008±0.0005 |
| Small cauc | 0.012±0.005 | 0.151±0.035 | 0.033±0.001 | 0.028±0.001 | 0.021±0.001 | 0.011±0.0001 | 0.011±0.0005 | 0.008±0.0005 |
| Large cauc | 0.282±0.065 | 3.148±0.880 | 0.033±0.001 | 0.029±0.001 | 0.021±0.001 | 0.011±0.0001 | 0.011±0.0001 | 0.008±0.0001 |
| Small fac | 0.017±0.011 | 2.025±1.854 | 0.029±0.001 | 0.029±0.001 | 0.022±0.001 | 0.010±0.0001 | 0.010±0.0001 | 0.008±0.0001 |
| Large fac | 0.732±0.324 | 6.229±2.672 | 0.030±0.001 | 0.031±0.001 | 0.022±0.001 | 0.010±0.0005 | 0.011±0.0001 | 0.008±0.0001 |

Table 13: Cost in time (in seconds) and GPU memory (in GBs) of an epoch of training, for our proposed tripartite MPNNs versus the bipartite ablation.

| | Tri. | MPNN | Small instances | | | | Large instances | | | |
|-----------|------|------|-----------------|--------|--------|--------|-----------------|--------|--------|--------|
| | | | Setcover | Indset | Cauc | Fac | Setcover | Indset | Cauc | Fac |
| T. (sec) | ✓ | GEN | 9.057 | 11.652 | 16.190 | 14.910 | 78.193 | 79.035 | 75.225 | 76.166 |
| | | GCN | 6.985 | 7.128 | 7.820 | 7.093 | 26.249 | 36.680 | 31.067 | 21.241 |
| | | GIN | 6.809 | 6.847 | 7.839 | 9.260 | 33.815 | 27.751 | 31.587 | 26.378 |
| | ✗ | GEN | 6.598 | 6.812 | 7.176 | 7.170 | 10.066 | 34.914 | 26.692 | 27.133 |
| | | GCN | 6.560 | 6.665 | 6.899 | 6.742 | 9.694 | 10.788 | 12.267 | 12.285 |
| | | GIN | 6.653 | 6.625 | 6.935 | 6.867 | 9.514 | 13.975 | 11.622 | 21.778 |
| Mem. (GB) | ✓ | GEN | 23.399 | 32.227 | 51.237 | 48.900 | 67.995 | 69.427 | 65.875 | 71.852 |
| | | GCN | 8.643 | 12.669 | 15.597 | 11.297 | 22.064 | 27.086 | 20.288 | 16.246 |
| | | GIN | 7.318 | 10.217 | 17.611 | 23.906 | 18.966 | 21.781 | 22.902 | 24.560 |
| | ✗ | GEN | 2.462 | 19.098 | 26.031 | 28.096 | 8.104 | 41.103 | 33.250 | 35.623 |
| | | GCN | 1.015 | 4.255 | 8.233 | 7.292 | 2.546 | 9.088 | 9.386 | 10.511 |
| | | GIN | 2.051 | 5.808 | 8.642 | 14.339 | 4.150 | 12.416 | 9.844 | 20.740 |