
Unsupervised Learning for Combinatorial Optimization with Principled Objective Design

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Using machine learning to solve combinatorial optimization (CO) problems is
2 challenging, especially when the data is unlabeled. This work proposes an unsu-
3 pervised learning framework for CO problems. Our framework follows a standard
4 relaxation-plus-rounding approach and adopts neural networks to parameterize
5 the relaxed solutions so that simple back-propagation can train the model end-to-
6 end. Our key contribution is the observation that if the relaxed objective satisfies
7 entry-wise concavity, a low optimization loss guarantees the quality of the final
8 integral solutions. This observation significantly broadens the applicability of a
9 previous framework inspired by Erdos’ probabilistic method [1]. In particular, this
10 observation can guide the design of objective models in the applications where
11 the objectives are not given explicitly while requiring being modeled in prior. We
12 evaluate our framework by solving a synthetic graph optimization problem, and
13 two real-world applications including resource allocation in circuit design and
14 approximate computing. Our framework largely outperforms the baselines based
15 on naïve relaxation, reinforcement learning and Gumbel-softmax tricks.

16 1 Introduction

17 Combinatorial optimization (CO) with the goal of finding the optimal solution from a discrete space is
18 a fundamental problem in many scientific and engineering applications [2–4]. Most CO problems are
19 NP-complete. Traditional methods efficient in practice often use heuristics or produce approximation
20 solutions. Designing these approaches requires considerable insights into the problem. Recently,
21 machine learning has revolutionized this traditional way to develop CO algorithms by leveraging
22 neural networks (NNs) to extract heuristics from the data [5–7]. Several learning for CO (LCO)
23 approaches have already been developed for SAT [8–10], mixed integer linear programming [11–13],
24 vertex covering [14, 15] and routing problems [16–23].

25 Another promising, if not more promising usage of machine learning techniques is to assist the
26 applications where the evaluation of the CO objective for each tentative solution could be expensive
27 and time-consuming [24–27]. For example, in hardware/system design, the actual computation
28 latency, power efficiency [28], and resource consumption [29–31] are unavailable before running
29 complex simulators. Also, in molecule design, the desired properties such as protein fluorescence or
30 DNA binding may only get evaluated via costly simulations or living experiments [32–34]. Therefore,
31 proxies of their objectives often need to be learned at first. And then, these Proxy-based CO (PCO)
32 can be solved further by following traditional LCO schemes [31]. Note that learning for PCO is even
33 in greater need compared to traditional CO problems because commercial CO solvers such as Gurobi
34 can never be applied in PCO due to the in-availability of closed-form objectives. Generic solvers
35 such as simulated annealing [35] may be applied while they could be extremely slow.

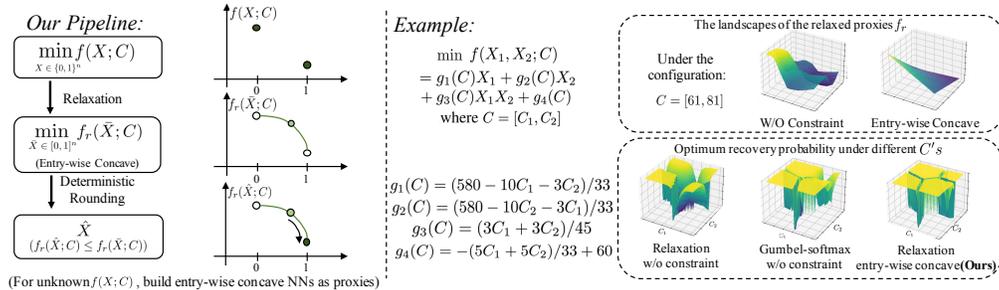


Figure 1: The pipeline and an example. The optimization objective $f(X; C)$, $X \in \{0, 1\}^n$ is relaxed to $f_r(\bar{X}; C)$, $\bar{X} \in [0, 1]^n$. Here, C is the problem configuration such as an attributed graph in a graph optimization problem. The entry-wise concave structure of f_r provides a performance guarantee in our deterministic rounding procedure (in Def. 1). When $f(\bar{X}; C)$ is not explicitly given, f_r will be modeled as NNs and learned. In the toy example, we first learn $f(\bar{X}; C)$ by proxies with or without the entry-wise concave constraint. We compare their landscapes in the top-right figure. We further optimize both proxies and round the obtained soft solutions to integral solutions. The bottom-right figure shows the optimum-recovery probabilities of different methods under different C 's.

36 In this work, we are to propose an unsupervised LCO framework. Our findings are applied to general
 37 CO problems while exhibiting extraordinary promise for PCO problems. Unsupervised LCO has
 38 recently attracted great attentions [1, 9, 20, 36, 37]. Compared to supervised learning that gets criticized
 39 for the dependence on huge amounts of labeled data [38], and reinforcement learning (RL) that
 40 suffers from notoriously unstable training [39], unsupervised LCO has shown great advantages due to
 41 its faster training, good generalization via accessing large unlabeled data, and its capability of dealing
 42 with large-scale problems [1]. Though with those advantages, unsupervised learning has never been
 43 investigated for PCO problems. Previous works for PCO problems, e.g., hardware design [30, 31],
 44 were all based on RL. This work provides the first unsupervised framework for PCO problems.

45 Our framework follows a relaxation-plus-rounding approach. We optimize a carefully-designed
 46 continuous relaxation of the cost model (penalized with constraints if any) and obtain a soft solution.
 47 Then, we decode the soft assignment to have the final discrete solution. This follows a common
 48 approach in traditional approximation algorithm design [40, 41]. However, the soft assignments
 49 here is given by a NN model optimized based on the historical (unlabeled) data via gradient descent.
 50 Learning from historical data is expected to facilitate the model understanding the data distribution,
 51 which helps with extracting heuristics, avoiding local minima and achieving fast inference. An
 52 illustration of the framework with a toy-example is shown in Fig. 1.

53 Our method shares a similar spirit with [1] while making the following significant contributions.
 54 We abandon the probabilistic guarantee in [1], because it is hard to use when we have general
 55 or proxy-based objectives. Instead, we design a deterministic objective relaxation principle that
 56 gives performance guarantee. We prove that if the objective relaxation is entry-wise concave w.r.t.
 57 the binary optimization variables, a low-cost soft solution plus deterministic sequential decoding
 58 guarantees generating a valid and low-cost integral solution. This principle significantly broadens the
 59 applicability of this unsupervised learning framework. In particular, it guides the design of model
 60 architectures to learn the objectives in PCO problems. We further justify the wide applicability of the
 61 entry-wise concave principle in both theory and practice.

62 We evaluate our framework over three PCO applications including feature-based edge covering & node
 63 matching problems, and two real-world applications, including imprecise functional unit assignment
 64 in approximate computing (AxC) [42–46] and resource allocation in circuit design [24, 30]. In all
 65 three applications, our framework achieves a significant performance boost compared to previous
 66 RL-based approaches and relaxed gradient approaches based on the Gumbel-softmax trick [47–49].

67 1.1 Further Discussion on Related Works

68 Most previous LCO approaches are based on RL [13, 14, 16–18, 22, 50–52] or supervised learning [11,
 69 12, 38], as these two frameworks do not hold much constraints on the formulation of CO problems.

70 However, they often suffer from the issues of training instability and subpar generalization. Previous
71 works on unsupervised learning for CO have studied satisfaction problems [9, 36], while applying
72 them to general CO problems requires problem reductions. Others have considered max-cut [37] and
73 TSP problems [20], while these works depend on carefully selected problem-specific objectives. The
74 work most relevant to ours is [1] and we give detailed comparison in Sec. 3. Note that all previous
75 works on unsupervised learning for CO do not apply to PCO as they need an explicit objective to
76 manipulate. For PCO problems, previous studies focus on how to learn more generalizable proxies of
77 the costs, such as via Bayesian learning [53, 54] and adversarial training [55, 56]. Once proxies are
78 learned, direct objective relaxation [55] or RL [30, 31] is often adopted. Studying generalization of
79 proxies is out of the scope of this work while entry-wise concave proxies seem smoother than those
80 without constraints (See Fig. 1) and thus have the potential to be more generalizable.

81 2 Preliminaries and Problem Formulation

82 In this section, we define several useful concepts and notations.

83 **Combinatorial Optimization (CO).** Let $C \in \mathcal{C}$ denote a data-based configuration such as a graph
84 with weighted edges. Let Ω be a finite set of all feasible combinatorial objects and each object has a
85 binary vector embedding $X = (X_i)_{1 \leq i \leq n} \in \{0, 1\}^n$. For example, in the node matching problem,
86 each entry of X corresponds to an edge to denote whether this edge is selected or not. Note that
87 such binary embeddings are applicable even when the choice is not naturally binary: Choosing at
88 most one element from a tuple $(1, 2, 3)$ can be represented as a 3-dim binary vector (X_1, X_2, X_3)
89 with the constraint $X_1 + X_2 + X_3 \leq 1$. W.l.o.g, we assume an algebraic form of the feasible set
90 $\Omega \triangleq \{X \in \{0, 1\}^n : g(X; C) < 1\}$ where $g(X; C) \geq 0$ for all $X \in \{0, 1\}^n$ ¹. For notational
91 simplicity, we only consider one inequality constraint while our later discussion in Sec. 3 and our
92 case studies in Sec. 4 may contain multiple inequalities. Given a configuration C and a constraint Ω ,
93 a combinatorial optimization (CO) is to minimize a cost $f(\cdot; C)$ by solving

$$\min_{X \in \{0, 1\}^n} f(X; C), \quad \text{s.t. } g(X; C) < 1. \quad (1)$$

94 **Proxy-based CO (PCO).** In the many applications, the cost or the constraint may not be cheaply
95 evaluated. Some proxies of the cost f or the constraint g often need to be learned from the historical
96 data. With some abuse of notations, we interchangeably use f (g , resp.) to denote the objective (the
97 constraint, resp.) and its proxy.

98 **Learning for CO/PCO (LCO).** A LCO problem is to learn an algorithm $\mathcal{A}_\theta(\cdot) : \mathcal{C} \rightarrow \{0, 1\}^n$, say
99 a neural network (NN) parameterized by θ to solve CO or PCO problems. Given a configuration
100 $C \in \mathcal{C}$, we expect \mathcal{A}_θ to (a) generate a valid solution $\hat{X} = \mathcal{A}_\theta(C) \in \Omega$ and (b) minimize $f(\hat{X}; C)$.

101 There are different approaches to learn \mathcal{A}_θ . Our focus is unsupervised learning approaches where
102 given a configuration C , the ground-truth solution X^* is not accessible during the training. θ can
103 only be optimized based on the knowledge of the cost and the constraint, or their proxies.

104 **Erdős' Probabilistic Method (EPM).** The EPM has recently been brought for LCO [1]. Specifi-
105 cally, The EPM formulates $\mathcal{A}_\theta(C)$ as a randomized algorithm that essentially gives a probabilistic
106 distribution over the solution space $\{0, 1\}^n$, which solves the optimization problem:

$$\min_{\theta} \mathbb{E}[l(X; C)], \quad \text{where } l(X; C) \triangleq f(X; C) + \beta 1_{g(X; C) \geq 1}, X \sim \mathcal{A}_\theta(C) \text{ and } \beta > 0. \quad (2)$$

107 Karaliyas & Loukas proved that with $\beta > \max_{X \in \Omega} f(X; C)$ and a small expected loss $\mathbb{E}[l(X; C)] <$
108 β , sampling a sufficiently large number of $\hat{X} \sim \mathcal{A}_\theta(C)$ guarantees the existence of a feasible $\hat{X} \in \Omega$
109 that achieves the cost $f(\hat{X}; C) \leq \mathbb{E}[l(X; C)]$ [1]. Although this guarantee makes EPM intriguing,
110 applying EPM in practice is non-trivial. We will explain the challenge in Sec. 3.1, which inspires our
111 solutions and further guides the objective design for general CO and PCO problems.

¹Normalization $(g(\cdot; C) - g_{\min}) / (g_{\min}^+ - g_{\min})$ where $g_{\min}^+ = \min_{X \in \{0, 1\}^n \setminus \Omega} g(X; C)$ and $g_{\min} = \min_{X \in \{0, 1\}^n} g(X; C)$ always satisfies the property. g_{\min}^+ , g_{\min} often can be easily estimated in practice.

112 3 The Relaxation Principle for Unsupervised LCO

113 In this section, we start with the practical issues of EPM. Then, we introduce our solutions by
 114 proposing a relaxation principle of the objectives, which gives performance guarantee for general
 115 practical unsupervised LCO.

116 3.1 Motivation: The Practical Issues of EPM

117 Applying EPM in practice has two fundamental difficulties. First, optimizing θ in Eq.(2) is generally
 118 hard as the gradient $\frac{dX}{d\theta}$ does not generally exist so the chain rule cannot be used. We discuss the
 119 potential solutions to this problem in Sec. 3.4. Second, EPM needs to sample a large number of
 120 $X \sim \mathcal{A}_\theta(C)$ for evaluation to achieve the performance guarantee in [1]. This is not acceptable where
 121 the evaluation per sample is time-consuming and expensive.

122 So, in practice, Karalias & Loukas consider a deterministic method. They view $\mathcal{A}_\theta(C) \in [0, 1]^n$ as
 123 the parameters of Bernouli distributions to generate the entries of X so $\mathbb{E}[X] = \mathcal{A}_\theta(C)$. First, they
 124 optimize $\min_\theta l(\mathcal{A}_\theta(C), C)$ instead of $\min_\theta \mathbb{E}[l(X, C)]$, and then, sequentially round the probability
 125 $\mathcal{A}_\theta(C)$ to discrete $X \in \{0, 1\}^n$ by comparing conditional expectations, e.g., $\mathbb{E}[l(X, C)|X_1 = 0]$ v.s.
 126 $\mathbb{E}[l(X, C)|X_1 = 1]$ to decide X_1 . However, such conditional expectations for general l cannot be
 127 efficiently computed unless one uses Monte-Carlo sampling. The two special cases in [1] on the
 128 max-clique and graph-partition problems seem to have special structures. This blocks the applicability
 129 of this framework, especially for the PCO problems where the objectives l are learned as models.

130 3.2 Our Approach: Relaxation plus Rounding, and Performance Guarantee

131 Our solution does not use the probabilistic modeling but directly adopts a relaxation-plus-rounding
 132 approach. We optimize a relaxation of the objective l_r and obtain a soft solution $\bar{X} \in [0, 1]^n$. Then,
 133 we deterministically round the entries in \bar{X} to a solution in the discrete space $\{0, 1\}^n$. The question
 134 is whether the obtained solutions may still achieve the guarantee as EPM does. Our key observation
 135 is that such success essentially depends on how to relax the objective l .

136 Therefore, our first contribution beyond [1] is to propose the principle (Def. 2) to relax general costs
 137 and constraints. With this principle, the unsupervised LCO framework can deterministically yield
 138 valid and low-cost solutions (Thm. 1) as the EPM guarantees, and is applied to any objective l .

139 First, we introduce the pipeline. Consider a relaxation of a deterministic upper bound of Eq.(2):

$$\min_{\theta} l_r(\theta; C) \triangleq f_r(\bar{X}; C) + \beta g_r(\bar{X}; C), \text{ where } \bar{X} = \mathcal{A}_\theta(C) \in [0, 1]^n, \beta > 0. \quad (3)$$

140 Here $f_r(\cdot; C) : [0, 1]^n \rightarrow \mathbb{R}$ is the relaxation of $f(\cdot; C)$, which satisfies $f_r(X; C) = f(X; C)$ for
 141 $X \in \{0, 1\}^n$. The relation between g_r and g is similar, i.e., $g_r(X; C) = g(X; C)$ for $X \in \{0, 1\}^n$.
 142 Here, we also use the fact that $g_r(X; C)$ provides a natural upper bound $1_{g(X; C) \geq 1} \leq g_r(X; C)$ for
 143 $X \in \{0, 1\}^n$ given the normalization of $g(X; C)$ adopted in Sec. 2.

144 Now, suppose the parameter θ gets optimized so that $l_r(\theta; C)$ is small. Further, we adopt the
 145 sequential rounding in Def. 1 to adjust the continuous solution $\bar{X} = \mathcal{A}_\theta(C)$ to discrete solution X .

146 **Definition 1** (Rounding). *Given a continuous vector $\bar{X} \in [0, 1]^n$ and an arbitrary order of the*
 147 *entries, w.o.l.g., $i = 1, 2, \dots, n$, round \bar{X}_i into 0 or 1 and fix all the other variables un-changed. Set*
 148 *$X_i = \arg \min_{j=0,1} f_r(X_1, \dots, X_{i-1}, j, \bar{X}_{i+1}, \dots, \bar{X}_n; C) + \beta g_r(X_1, \dots, X_{i-1}, j, \bar{X}_{i+1}, \dots, \bar{X}_n; C)$,*
 149 *replace \bar{X}_i with X_i and repeat the above procedure until all the variables become discrete.*

150 Note that our rounding procedure does not need to evaluate any conditional expectations
 151 $\mathbb{E}[l(X; C)|X_i]$ which EPM in [1] requires. Instead, we ask both relaxations f_r and g_r to sat-
 152 isfy the principle in Def. 2. With this principle, the pipeline allows achieving a valid and low-cost
 153 solution X , as proved in Theorem 1. We leave the proof in Appendix A.1.

154 **Definition 2** (The Entry-wise Concave Principle). *For any $C \in \mathcal{C}$, $h_r(\cdot; C) : [0, 1]^n \rightarrow \mathbb{R}$ is*
 155 *entry-wise concave if for any $\gamma \in [0, 1]$ and any $\bar{X}, \bar{X}' \in [0, 1]^n$ that are only different in one entry,*

$$\gamma h_r(\bar{X}; C) + (1 - \gamma) h_r(\bar{X}'; C) \leq h_r(\gamma \bar{X} + (1 - \gamma) \bar{X}'; C).$$

156 Note that entry-wise concavity is much weaker than concavity. For example, the function
 157 $h_r(\bar{X}_1, \bar{X}_2) = -\text{Relu}(\bar{X}_1 \bar{X}_2)$, $\bar{X}_1, \bar{X}_2 \in \mathbb{R}$ is entry-wise concave but not concave.

158 **Theorem 1** (Performance Guarantee). *Let $\beta > \max_{X \in \Omega} f(X; C)$ and $\min_{X \in \Omega} f(X; C) \geq 0$ in
 159 Eq.(3). Suppose the relaxed cost f_r and constraint g_r are entry-wise concave, and the learned
 160 parameter θ achieves $l_r(\theta; C) < \beta$. Then, rounding (Def. 1) the relaxed solution $\bar{X} = \mathcal{A}_\theta(C)$
 161 generates a valid discrete solution $X \in \Omega$ such that $f(X; C) < l_r(\theta; C)$.*

162 When there are multiple normalization constraints $g^{(j)}(X; C) < 1$ for $j = 1, 2, \dots$, we may use
 163 relaxation $\beta \sum_j g_r^{(j)}(X; C)$ as the penalty term in Eq.(3), where $g_r^{(j)}$ is a relaxation of $g^{(j)}$. It can be
 164 shown that if $\sum_j g_r^{(j)}$ satisfies the entry-wise concave condition, the guarantee of Thm. 1 still applies.

165 3.3 The Wide Applicability of Entry-wise Concave Relaxations

166 We have introduced the entry-wise concave principle to relax the objective to associate our framework
 167 with performance guarantee. The question is how widely applicable this principle could be.

168 Actually, every function with binary inputs can be relaxed as an entry-wise affine function with the
 169 exactly same values at the discrete inputs, as shown in Theorem 2. Note that entry-wise affinity is a
 170 special case of entry-wise concavity. In Sec. 4, we will provide the design of NN architecture (for
 171 PCO) and math derivation (for CO) that guarantee formulating an entry-wise concave function.

172 **Theorem 2** (Wide Applicability). *For any binary-input function $h(\cdot) : \{0, 1\}^n \rightarrow \mathbb{R}$, there exists a
 173 relaxation $h_r(\cdot) : [0, 1]^n \rightarrow \mathbb{R}$ such that (a) $h_r(X) = h(X)$ for $X \in \{0, 1\}^n$ and (b) h_r is entry-wise
 174 affine, i.e., for any $\gamma \in [0, 1]$ and any $\bar{X}, \bar{X}' \in [0, 1]^n$ that are only different in one entry,*

$$\gamma h_r(\bar{X}) + (1 - \gamma) h_r(\bar{X}') = h_r(\gamma \bar{X} + (1 - \gamma) \bar{X}').$$

175 *Proof sketch.* Set $h_r(\bar{X}) = \sum_{X \in \{0, 1\}^n} h(X) \prod_{j=1}^n \bar{X}_j^{X_j} (1 - \bar{X}_j)^{(1 - X_j)}$, which satisfies (a) and (b).
 176 Note that we suppose that $\bar{X}_j^0 = 1$ for any $\bar{X}_j \in [0, 1]$. The detailed proof is in Appendix A.2. \square

177 Although Theorem 2 shows the existence of entry-wise affine relaxations, the constructed representa-
 178 tion in the proof depends on higher-order moments of the input entries, which make it often intractable
 179 to implement, especially via a NN architecture. Therefore, we also propose to use entry-wise concave
 180 functions to implicitly generate higher-order moments. For example, when $n = 2$, we could use
 181 the composition of $-\text{Relu}(\cdot)$ and affine operators (only 1st-order moments) to achieve universal
 182 representation (See Prop. 1 and the proof in Appendix A.3). For general n , we leave as a future study.

183 **Proposition 1.** *For any binary-input function $h(X_1, X_2)$, there exists parameters $\{w_{ij}\}$ such that
 184 an entry-wise concave function $h_r(\bar{X}_1, \bar{X}_2) = w_{00} - \sum_{i=1}^3 \text{Relu}(w_{i1} \bar{X}_1 + w_{i2} \bar{X}_2 + w_{i0})$ satisfies
 185 $h_r(X_1, X_2) = h(X_1, X_2)$ for any $X_1, X_2 \in \{0, 1\}$.*

186 3.4 Discussion: Methods to Directly Optimize the Randomized Objective in EPM Eq.(2)

187 The naïve way to optimize the randomized
 188 objective in Eq.(2) without worrying about
 189 the specific form of the objective l is based
 190 on the policy gradient in RL via the loga-
 191 rithmic trick, i.e., estimating the gradient $\frac{dl}{d\theta}$
 192 via $(f(X; C) + \beta 1_{g(X; C) \geq 1}) \log \mathbb{P}(X)$ by sam-
 193 pling $X \sim \mathcal{A}_\theta(C)$. However, the policy gradi-
 194 ent suffers from notoriously large variance [39]
 195 and makes RL hard to converge. Therefore,
 196 methods such as actor critic [57] or subtracting
 197 some baselines $l(X; C) - b$ [58] have been proposed.

	RL	Gumbel-softmax	Ours
Objective	No Limit	No Limit	Entry-wise Concave
Optimizer	Log Trick	Gumbel Trick	No Limit
Inference	Sampling	Sampling	Deter. Rounding
Train. Time	Slow	Fast	Fast
Convergence	Hard	Medium	Easy
Infer. Time	Slow	Slow	Fast

Table 1: The comparison among RL (policy gradient), Gumbel-softmax methods and our principled objective relaxation. Our methods are in need of much less training time and inference time.

198 Another way to solve Eq.(2) is based on reparameterization tricks to reduce the variance of gradi-
 199 ents [59, 60]. Specifically, we set the entries of output $\bar{X} = \mathcal{A}_\theta(C) \in [0, 1]^n$ as the parameters
 200 of Bernoulli distributions to generate X , i.e., $X_i \sim \text{Bern}(\bar{X}_i)$, for $1 \leq i \leq n$. To make $dX_i/d\bar{X}_i$

201 computable, we may use the Gumble-softmax trick [47–49]. However, this approach suffers from two
 202 issues. First, the estimation of the gradient is biased. Second, as $\mathcal{A}_\theta(C)$ is essentially a randomized
 203 algorithm, sampling sufficiently many $X \sim \mathcal{A}_\theta(C)$ is needed to guarantee a valid and low-cost
 204 solution. However, such evaluation is costly as discussed in Sec. 3.1. We compare different aspects
 205 of RL, Gumbel-softmax tricks and our relaxation approach in Table 1.

206 4 Applying Our Relaxation Principle to Learning for PCO

207 In this section, we apply our relaxation principle to three PCO applications: (I) feature-based edge
 208 covering & node matching, (II) resource allocation in circuit design, and (III) imprecise functional
 209 unit assignment in approximate computing. All the applications have graph-based configurations C .
 210 So later, we first introduce how to use graph neural networks (GNNs) to build proxies that satisfy
 211 our relaxation principle. Such GNN-based proxies will be used as the cost function relaxation f_r
 212 in all the applications. Our principle can also guide the relaxation of explicit CO objectives. The
 213 constraints in applications (I)(III) are explicit and their relaxation can be written into the entry-wise
 214 affine form. The constraint in (II) needs another GNN-based entry-wise concave proxy to learn.

215 4.1 GNN-based Entry-wise Concave Proxies

216 We consider the data configuration C as an attributed graph (V, E, Z) where V is the node set,
 217 $E \subseteq V \times V$ is the edge set and Z is the node attributes. We associate each node with a binary variable
 218 and group them together $X : \in \{0, 1\}^{|V|}$, where for each $v \in V$, $X_v = 1$ indicates the choice of the
 219 node v . Note that our approach can be similarly applied to edge-level variables (see Appendix C.2),
 220 which is used in application (I). Let \bar{X} still denote the relaxation of X .

221 To learn a discrete function $h : \{0, 1\}^{|V|} \times \mathcal{C} \rightarrow \mathbb{R}$, we adopt a GNN as the relaxed proxy of h . We
 222 first define a latent graph representation in \mathbb{R}^F whose entries are all entry-wise affine mappings of X .

$$\text{Latent representation: } \phi(\bar{X}; C) = W + \sum_{v \in V} U_v \bar{X}_1 + \sum_{v, u \in V, (v, u) \in E} Q_{v, u} \bar{X}_v \bar{X}_u \quad (4)$$

223 where W is the graph representation, U_v 's are node representations and $Q_{v, u}$ are edge representations.
 224 These representations do not contain X and are given by GNN encoding C . Here, we consider at
 225 most 2nd-order moments based on adjacent nodes as they can be easily implemented via current
 226 GNN platforms [61, 62]. Then, we use ϕ to generate entry-wise affine & concave proxies as follows.

$$\text{Entry-wise Affine Proxy (AFF): } h_r^a(\bar{X}; C) = \langle w^a, \phi(\bar{X}; C) \rangle. \quad (5)$$

$$\text{Entry-wise Concave Proxy (CON): } h_r^c(\bar{X}; C) = \langle w^c, -\text{Relu}(\phi(\bar{X}; C)) \rangle + b. \quad (6)$$

227 where $w^a, w^c \in \mathbb{R}^F, b \in \mathbb{R}$ are learnt parameters and $w^c \geq 0$ guarantees entry-wise concavity.

228 4.2 The Setting up of the Experiments

229 **Training & Evaluation Pipeline.** In all the applications, we
 230 adopt the following training & evaluation pipeline. First, we
 231 have a set of observed configurations $\mathcal{D}_1 \subset \mathcal{C}$. Each $C \in \mathcal{D}_1$ is
 232 paired with one $X \in \{0, 1\}^n$. We use the costs $f(X, C)$ (and
 233 constraints $g(X, C)$) to train the relaxed proxies $f_r(X, C)$ (and
 234 $g_r(X, C)$, if cannot be derived explicitly), where the relaxed proxies follow either Eq.(5) (named
 235 AFF) or Eq.(6) (named CON). Then, we parameterize the LCO algorithm $\mathcal{A}_\theta(C) \in [0, 1]^n$ via another
 236 GNN. Based on the learned (or derived) f_r and g_r , we optimize θ by minimizing $\sum_{C \in \mathcal{D}_1} l_r(\theta; C)$,
 237 where l_r is defined according to Eq.(3). We will split \mathcal{D}_1 into a training set and a validation set for
 238 hyperparameter-tuning of proxies and \mathcal{A}_θ . We have another set of configurations $\mathcal{D}_2 \subset \mathcal{C}$ used for
 239 testing. For each $C \in \mathcal{D}_2$, we use the relaxation $\bar{X} = \mathcal{A}_\theta(C)$ plus our rounding to evaluate the
 240 learned algorithm $\mathcal{A}_\theta(\cdot)$. We follow [1] and do not consider fine-tuning \mathcal{A}_θ over the testing dataset
 241 \mathcal{D}_2 to match the potential requirement of the fast inference.

242 **Baselines.** We consider 4 common baselines that is made up of different learnable relaxed proxies
 243 f_r, g_r , algorithms \mathcal{A}_θ and inference approaches as shown in Table 2. For the proxies f_r, g_r for

Baseline	f_r, g_r	\mathcal{A}_θ	Inference
Naïve + R	no limit	no limit	rounding
RL	no limit	RL	sampling
GS-Tr+S	no limit	GS	sampling
GS-Tr+R	no limit	GS	rounding

Table 2: The baselines in the paper.

244 baselines, we apply GNNs without the entry-wise concave constraint and use X as one node attribute
 245 while keeping all other hyper-parameters exactly the same as ours (See details in Appendix. C); For the
 246 algorithm \mathcal{A}_θ , we provide the Gumbel-softmax trick based methods (GS-Tr) [48, 49], the actor-critic-
 247 based RL method [57] (RL) and the naïve relaxation method (Naïve); For the inference approaches,
 248 we consider Monte Carlo sampling (S) and our proposed rounding (R) procedure. Although the
 249 baselines adopt proxies that are different from ours, we guarantee that their proxies approximate
 250 the ground-truth f, g over the validation dataset at least no worse than ours. In application II, we
 251 also consider two non-learnable algorithms to optimize the proxies without relaxation constraints,
 252 simulated annealing (SA) [35] and genetic algorithms (GA) [63, 64]. In application III, we put
 253 all of the required AxC units either close to the input (C-In) or close to the output (C-Out) of the
 254 approximating computing circuit as additional baselines. More details of the experiments setups and
 255 hyperparameter tuning can be found in Appendix C. We also obtain the optimal solutions (OPT) for
 256 applications I and III via brute-force search for comparison.

257 4.3 Application I: Feature-based Edge Covering & Node Matching in Graphs

258 This application is inspired by [65]. Here, each configuration C is a 4×4 grid graph whose node
 259 attributes are two-digit images generated by random combinations of the pictures in MNIST [66].
 260 We associated each edge with variables $X \in \{0, 1\}^{|E|}$. The objective is the sum of edge weights
 261 $f(X; C) = \sum_{e \in E} w_e X_e$ where w_e is unknown in prior and needed to be learned. The ground truth
 262 of w_e is a multiplication of the numbers indicated by the images on the two adjacent nodes. We adopt
 263 ResNet-50 [67] (to refine node features) plus GraphSAGE [68] to encode C . We consider using both
 264 Eq.(5) and Eq.(6) to formulate the relaxed cost $f_r(\bar{X}; C)$. Training and validating f_r are based on
 265 100k randomly sampled C paired with randomly sampled X . Note that 100k is much smaller than
 266 the entire space $\{0, 1\}^{|E|} \times \mathcal{C}$ is of size $2^{24} \times 100^{16}$.

267 Next, as the constraint here is explicit, we can derive the relaxation of the constraints for this
 268 application. First, the constraint relaxation of the edge covering problem can be written as

$$\text{Edge Covering Constraint: } g_r(\bar{X}; C) = \sum_{v \in V} \prod_{e: v \in e} (1 - \bar{X}_e). \quad (7)$$

269 Each production term in Eq.(7) indicates that for each node, at least one edge is selected. We can
 270 easily justify that g_r is entry-wise affine and $\Omega = \{X \in \{0, 1\}^{|E|} : g_r(X; C) < 1\}$ exactly gives the
 271 feasible solutions to the edge covering problem.

272 Similarly, we can derive the constraint for node matching by adding a further term to Eq.(7).

$$\text{Node Matching Constraint: } g_r(\bar{X}; C) = \sum_{v \in V} \left[\prod_{e: v \in e} (1 - \bar{X}_e) + \prod_{e_1, e_2: v \in e_1, e_2, e_1 \neq e_2} \bar{X}_{e_1} \bar{X}_{e_2} \right]. \quad (8)$$

273 Here, the second term indicates that no two edges adjacent to the same node can be selected. We can
 274 easily justify that g_r is entry-wise affine and $\Omega = \{X \in \{0, 1\}^{|E|} : g_r(X; C) < 1\}$ groups exactly
 275 the feasible solutions to the node matching problem.

276 Note that our above derivation also generalizes the node-selection framework in [1] to edge selection.
 277 With the learned f_r and the derived g_r , we further train and validate \mathcal{A}_θ over the 100k sampled
 278 (X, C) 's and test on another 500 randomly sampled C 's.

279 **Evaluation.** Table 3 shows the evaluation results. In the
 280 GS-Tr+S method, the number of sampling is set to 120
 281 (about 2.5 times the inference time of our deterministic
 282 rounding). Note that for node matching, GS-Tr+S could
 283 hardly sample a feasible node matching solution within
 284 120 samples. The experiment results show that our
 285 principled proxy design exceeds the other baselines on
 286 both tasks. Also, we observe that AFF outperforms
 287 CON, which results from the fact that $f(X; C)$ in these
 288 two problems are naturally in entry-wise affine forms

Method	Edge covering	Node matching
Naive+R	68.52	429.12
RL	51.29	426.97
GS-Tr+S	63.36	-
GS-Tr+R	46.91	429.39
CON(ours)	49.59	422.47
AFF(ours)	44.55	418.96
OPT(gt)	42.69	416.01

Table 3: Performance on application I (graph optimization).

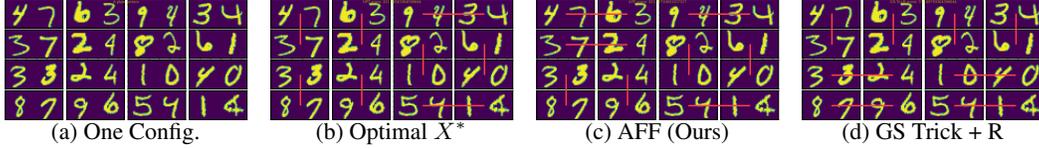


Figure 2: The visualization for node matching in Application I. Our method avoids large multiplications $87 * 96$ and $94 * 82$ where GS-Trick cannot, and generate a solution different but close to OPT.

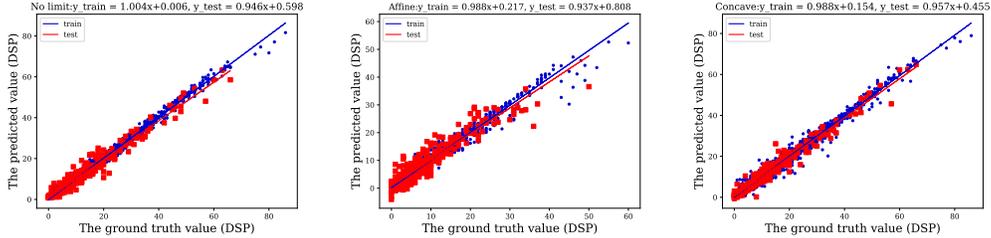


Figure 3: Comparing different proxies for learning DSP usage. Left, no constraint; Middle, entry-wise affine constraint (Eq. (5)); Right, entry-wise concave constraint (Eq.(6))

289 with low-order (1st-order) moments. One instance of node matching randomly selected from the test
 290 set is shown in Fig. 2. More visualization results can be found in Fig. 5 in the appendix.

291 4.4 Application II: Resource Allocation in Circuit Design

292 Resource allocation in field-programmable gate array (FPGA) design is a fundamental problem which
 293 can lead to largely varied circuit quality after synthesis, such as area, timing, and latency. In this
 294 application, we follow the problem formulation in [24, 30], where the circuit is represented as a data
 295 flow graph (DFG), and each node represents an arithmetic operation such as multiplication or addition.
 296 The goal is to find a resource allocation for each node to be either digital signal processor (DSP) or
 297 look-up table (LUT), such that the final circuit area (i.e., actual DSP and LUT usage) after synthesis
 298 is minimized. Notably, different allocation solutions result in greatly varied DSP/LUT usage due to
 299 complicated synthesis process, which cannot be simply summed up over each node. To obtain precise
 300 DSP/LUT usage, one must run high-level synthesis (HLS) [69] and place-and-route [70] tools, which
 301 can take up to hours [24, 30].

302 In this application, each configuration C is a DFG with > 100 nodes, where each node is allocated
 303 to either DSP or LUT. Node attributes include operation type (i.e., multiplication or addition) and
 304 data bitwidth. More details about the dataset can be found in Appendix C.4. Let $X \in \{0, 1\}^{|V|}$
 305 denote the mapping to DSP or LUT. Let f_r and g_r denote the proxies of actual LUT and actual DSP
 306 usage, respectively. Note that given different constraints on the DSP usage, we will normalize g_r as
 307 introduced in Sec. 2. We train and validate $f_r, g_r, \mathcal{A}_\theta$ on 8,000 instances that consist of 40 DFGs
 308 (C), each DFG with 200 different mappings (X), and test \mathcal{A}_θ over 20 DFGs. Note that the actual
 309 LUT and DSP usages of each training instance has been collected by running HLS in prior. We also
 310 run HLS to evaluate the actual LUT and DSP usages for the testing cases given the learned mappings.

311 **Evaluation.** We rank each method’s best actual LUT usage under the constraint of different percent-
 312 ages (40% - 70%) of the maximum DSP usage in each testing instance, then calculate the averaged
 313 ranks. Fig. 4 shows the results. Our entry-wise concave proxy achieves the best performance.
 314 GS-Tr+R is slightly better than RL, and both of them exceed SA and GA. We do not include our
 315 entry-wise affine proxy in the ranking list, because the affine proxy could be much less accurate
 316 than the proxy without constraints and the entry-wise concave proxy. The comparison between
 317 these proxies on learning DSP usage (& LUT usage) is shown in Fig. 3 (& Fig. 7 in the appendix,
 318 respectively). The gap between different proxies indicates the FPGA circuit contains high-order
 319 moments of the input optimization variables and 2-order entry-wise affine proxy cannot model well.
 320 We do not include the result of GS-Tr+S and Naive+R, because these methods perform poor and
 321 could hardly generate feasible solutions given a constraint of DSP usage. We leave their results in
 322 Table. 6 in the appendix. Moreover, we compare the training time between different methods. To

DSP usage	40%	45%	50%	55%	60%	65%	70%	rank-avg
SA	3.50	3.25	3.42	3.17	3.50	4.08	4.00	3.56
GA	2.70	2.92	3.17	3.08	3.42	3.25	3.25	3.11
RL	3.20	3.67	3.67	3.17	2.83	2.58	2.33	3.06
GS-Tr+R	3.50	3.00	2.50	3.08	2.17	2.50	2.75	2.79
CON	2.10	2.17	2.25	2.25	3.00	2.50	2.50	2.40

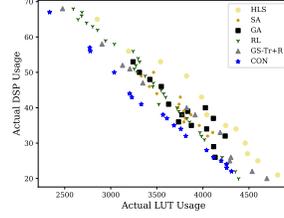


Figure 4: The left table shows averaged ranks of the LUT usage given by different methods with the constraint of different percentage of DSP usage in Application II (resource allocation). The right figure shows the DSP-LUT usage amount relationship on one test configuration. The HLS baseline denotes the optimal synthesis results among 200 random mappings.

Threshold θ	C-In	C-Out	Naïve	RL	GS-Tr+S	GS-Tr+R	CON	AFF	OPT
3 AxC units	12.42	12.44	3.62	10.59	4.87	3.24	3.18	3.10	2.77
5 AxC units	14.68	14.65	6.20	12.28	8.03	5.86	5.13	5.38	4.74
8 AxC units	17.07	17.04	11.12	15.17	12.65	10.62	10.17	10.04	8.56

Table 4: Relative errors of different methods with the AxC unit constraint as 3,5,8 in Application III.

323 be fair, all methods run on the same server with a Quadro RTX 6000 GPU. The RL based optimizer
 324 takes 22 GB GRAM, while other optimizers only take 7 GB on average. Fig. 8 in the appendix further
 325 demonstrates that our methods and GS-T methods require much less training time than RL.

326 4.5 Application III: Imprecise Functional Unit Assignment in Approximate Computing

327 One fundamental problem in approximate computing (AxC) is to assign imprecise functional units
 328 (a.k.a., AxC units) to execute operations such as multiplication or addition [42–46], aiming to
 329 significantly reduce circuit energy with tolerable error. We follow the problem formulation in [45],
 330 where given a computation graph, each node represents either multiplication or addition. The
 331 incoming edges of a node represent its two operands. The goal is to assign AxC units to a certain
 332 number of nodes while minimizing the expected relative error of the output of the computation graph.

333 In this application, each configuration C is a computation graph with 15 nodes (either multiplication
 334 or addition) that maps a vector in \mathbb{R}^{16} to \mathbb{R} . A fixed number θ of nodes are assigned to AxC units
 335 with produce 10% relative error. Let $X \in \{0, 1\}^{|V|}$ denote whether a node is assigned to an AxC unit
 336 or not; the proxy of the objective f_r is the expected relative error at the output. We use 100k (X, C)
 337 as the training dataset and the entire solution space is $2^{15} \times 2^{15}$. For each (X, C) , the ground-truth,
 338 i.e., expected relative error, is computed by averaging 1k inputs sampled uniformly at random from
 339 $[0, 10]^{16}$. The constraint g_r is $\sum_{v \in V} X_v \geq \theta$ with normalization, where $\theta \in \{3, 5, 8\}$. We test the
 340 learned \mathcal{A}_θ on 500 unseen configurations.

341 **Evaluation.** Table. 4 shows the averaged relative errors of the assignments by different methods. The
 342 problem is far from trivial. Intuitively, assigning AxC units closed to the output, we may expect small
 343 error. However, C-Out performs bad. Our proxies AFF and CON obtain comparable best results.
 344 The MAE loss values of the two proxies are also similar, as shown in Table 8 in the appendix. The
 345 reason is that the circuit is made up of 4 layers in total which leads to at most 4-order moments in
 346 the objective function, which is in a medium-level complexity. Training time is also studied for this
 347 application, resulting in the same conclusion as application II (See Table 7 in the appendix).

348 5 Conclusion

349 This work introduces an unsupervised end-to-end framework to resolve LCO problems based on
 350 the relaxation-plus-rounding technique. With our entry-wise concave architecture, our framework
 351 guarantees that a low objective value could lead to qualified discrete solutions. Our framework is
 352 particularly good at solving PCO problems where the objectives need to be modeled and learned.
 353 Real-world applications demonstrate the superiority of our method over RL and gradient-relaxation
 354 approaches in both optimization performance and training efficiency. In the future, we aim to further
 355 broaden our framework so that binary embeddings of the optimization variables are not needed.

References

- 356
- 357 [1] N. Karalias and A. Loukas, “Erdos goes neural: an unsupervised learning framework for
358 combinatorial optimization on graphs,” *Advances in Neural Information Processing Systems*,
359 vol. 33, 2020.
- 360 [2] C. H. Papadimitriou and K. Steiglitz, *Combinatorial optimization: algorithms and complexity*.
361 Courier Corporation, 1998.
- 362 [3] G. Naseri and M. A. Koffas, “Application of combinatorial optimization strategies in synthetic
363 biology,” *Nature communications*, vol. 11, no. 1, 2020.
- 364 [4] Y. Crama, “Combinatorial optimization models for production scheduling in automated manu-
365 facturing systems,” *European Journal of Operational Research*, vol. 99, no. 1, 1997.
- 366 [5] J. J. Hopfield and D. W. Tank, ““neural” computation of decisions in optimization problems,”
367 *Biological Cybernetics*, vol. 52, no. 3, 1985.
- 368 [6] K. A. Smith, “Neural networks for combinatorial optimization: a review of more than a decade
369 of research,” *INFORMS Journal on Computing*, vol. 11, no. 1, 1999.
- 370 [7] O. Vinyals, M. Fortunato, and N. Jaitly, “Pointer networks,” *Advances in Neural Information
371 Processing Systems*, vol. 28, 2015.
- 372 [8] D. Selsam, M. Lamm, B. Benedikt, P. Liang, L. de Moura, D. L. Dill *et al.*, “Learning a sat
373 solver from single-bit supervision,” in *International Conference on Learning Representations*,
374 2018.
- 375 [9] S. Amizadeh, S. Matushevych, and M. Weimer, “Learning to solve circuit-sat: An unsupervised
376 differentiable approach,” in *International Conference on Learning Representations*, 2018.
- 377 [10] E. Yolcu and B. Póczos, “Learning local search heuristics for boolean satisfiability,” *Advances
378 in Neural Information Processing Systems*, vol. 32, 2019.
- 379 [11] E. Khalil, P. Le Bodic, L. Song, G. Nemhauser, and B. Dilkina, “Learning to branch in mixed
380 integer programming,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 30,
381 no. 1, 2016.
- 382 [12] M. Gasse, D. Chételat, N. Ferroni, L. Charlin, and A. Lodi, “Exact combinatorial optimization
383 with graph convolutional neural networks,” *Advances in Neural Information Processing Systems*,
384 vol. 32, 2019.
- 385 [13] A. Delarue, R. Anderson, and C. Tjandraatmadja, “Reinforcement learning with combinatorial
386 actions: An application to vehicle routing,” *Advances in Neural Information Processing Systems*,
387 vol. 33, 2020.
- 388 [14] E. Khalil, H. Dai, Y. Zhang, B. Dilkina, and L. Song, “Learning combinatorial optimization
389 algorithms over graphs,” *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- 390 [15] Z. Li, Q. Chen, and V. Koltun, “Combinatorial optimization with graph convolutional networks
391 and guided tree search,” *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- 392 [16] I. Bello, H. Pham, Q. V. Le, M. Norouzi, and S. Bengio, “Neural combinatorial optimization with
393 reinforcement learning,” *International Conference on Learning Representations (Workshop)*,
394 2017.
- 395 [17] X. Chen and Y. Tian, “Learning to perform local rewriting for combinatorial optimization,”
396 *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- 397 [18] W. Kool, H. van Hoof, and M. Welling, “Attention, learn to solve routing problems!” in
398 *International Conference on Learning Representations*, 2018.
- 399 [19] C. K. Joshi, T. Laurent, and X. Bresson, “An efficient graph convolutional network technique
400 for the travelling salesman problem,” *arXiv preprint arXiv:1906.01227*, 2019.
- 401 [20] B. Hudson, Q. Li, M. Malencia, and A. Prorok, “Graph neural network guided local search
402 for the traveling salesperson problem,” *International Conference on Learning Representations*,
403 2022.

- 404 [21] Y. Ma, J. Li, Z. Cao, W. Song, L. Zhang, Z. Chen, and J. Tang, "Learning to iteratively solve
405 routing problems with dual-aspect collaborative transformer," *Advances in Neural Information
406 Processing Systems*, vol. 34, 2021.
- 407 [22] Y.-D. Kwon, J. Choo, I. Yoon, M. Park, D. Park, and Y. Gwon, "Matrix encoding networks
408 for neural combinatorial optimization," *Advances in Neural Information Processing Systems*,
409 vol. 34, 2021.
- 410 [23] M. Kim, J. Park *et al.*, "Learning collaborative policies to solve np-hard routing problems,"
411 *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- 412 [24] N. Wu, H. Yang, Y. Xie, p. Li, and C. Hao, "High-level synthesis performance prediction
413 using gnns: Benchmarking, modeling, and advancing," in *Proceedings of IEEE/ACM Design
414 Automation Conference (DAC), 2022.*, 2022.
- 415 [25] C. Mendis, A. Renda, S. Amarasinghe, and M. Carbin, "Ithemal: Accurate, portable and fast
416 basic block throughput estimation using deep neural networks," in *International Conference on
417 Machine Learning*, 2019.
- 418 [26] G. Zhang, H. He, and D. Katabi, "Circuit-gnn: Graph neural networks for distributed circuit
419 design," in *International Conference on Machine Learning*. PMLR, 2019.
- 420 [27] S. Vasudevan, W. J. Jiang, D. Bieber, R. Singh, C. R. Ho, C. Sutton *et al.*, "Learning semantic
421 representations to verify hardware designs," *Advances in Neural Information Processing Systems*,
422 vol. 34, 2021.
- 423 [28] N. Mishra, C. Imes, J. D. Lafferty, and H. Hoffmann, "Caloree: Learning control for predictable
424 latency and low energy," in *Proceedings of the Twenty-Third International Conference on
425 Architectural Support for Programming Languages and Operating Systems*, 2018.
- 426 [29] A. Renda, Y. Chen, C. Mendis, and M. Carbin, "Diff tune: Optimizing cpu simulator param-
427 eters with learned differentiable surrogates," in *2020 53rd Annual IEEE/ACM International
428 Symposium on Microarchitecture (MICRO)*. IEEE, 2020.
- 429 [30] N. Wu, Y. Xie, and C. Hao, "Ironman: Gnn-assisted design space exploration in high-level
430 synthesis via reinforcement learning," in *Proceedings of the 2021 on Great Lakes Symposium
431 on VLSI*, 2021.
- 432 [31] A. Mirhoseini, A. Goldie, M. Yazgan, J. W. Jiang, E. Songhori, S. Wang, Y.-J. Lee, E. Johnson,
433 O. Pathak, A. Nazi *et al.*, "A graph placement methodology for fast chip design," *Nature*, vol.
434 594, no. 7862, 2021.
- 435 [32] K. Chen and F. H. Arnold, "Enzyme engineering for nonaqueous solvents: random mutagenesis
436 to enhance activity of subtilisin e in polar organic media," *Bio/Technology*, vol. 9, no. 11, 1991.
- 437 [33] C. Angermueller, D. Dohan, D. Belanger, R. Deshpande, K. Murphy, and L. Colwell, "Model-
438 based reinforcement learning for biological sequence design," in *International Conference on
439 Learning Representations*, 2019.
- 440 [34] R. Gómez-Bombarelli, J. N. Wei, D. Duvenaud, J. M. Hernández-Lobato, B. Sánchez-Lengeling,
441 D. Sheberla, J. Aguilera-Iparraguirre, T. D. Hirzel, R. P. Adams, and A. Aspuru-Guzik, "Au-
442 tomatic chemical design using a data-driven continuous representation of molecules," *ACS
443 Central Science*, vol. 4, no. 2, 2018.
- 444 [35] D. Bertsimas and J. Tsitsiklis, "Simulated annealing," *Statistical Science*, vol. 8, no. 1, 1993.
- 445 [36] J. Toenshoff, M. Ritzert, H. Wolf, and M. Grohe, "Run-csp: unsupervised learning of message
446 passing networks for binary constraint satisfaction problems," 2019.
- 447 [37] W. Yao, A. S. Bandeira, and S. Villar, "Experimental performance of graph neural networks
448 on random instances of max-cut," in *Wavelets and Sparsity XVIII*, vol. 11138. International
449 Society for Optics and Photonics, 2019.
- 450 [38] G. Yehuda, M. Gabel, and A. Schuster, "It's not what machines can learn, it's what we cannot
451 teach," in *International Conference on Machine Learning*. PMLR, 2020.

- 452 [39] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves,
453 M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, “Human-level control through deep rein-
454 forcement learning,” *Nature*, vol. 518, no. 7540, 2015.
- 455 [40] R. Gandhi, S. Khuller, S. Parthasarathy, and A. Srinivasan, “Dependent rounding and its
456 applications to approximation algorithms,” *Journal of the ACM (JACM)*, vol. 53, no. 3, 2006.
- 457 [41] J. Byrka, F. Grandoni, T. Rothvoß, and L. Sanità, “Steiner tree approximation via iterative
458 randomized rounding,” *Journal of the ACM (JACM)*, vol. 60, no. 1, 2013.
- 459 [42] J. Han and M. Orshansky, “Approximate computing: An emerging paradigm for energy-efficient
460 design,” in *2013 18th IEEE European Test Symposium (ETS)*. IEEE, 2013.
- 461 [43] S. Mittal, “A survey of techniques for approximate computing,” *ACM Computing Surveys*
462 (*CSUR*), vol. 48, no. 4, 2016.
- 463 [44] R. Venkatesan, A. Agarwal, K. Roy, and A. Raghunathan, “Macaco: Modeling and analysis
464 of circuits for approximate computing,” in *2011 IEEE/ACM International Conference on*
465 *Computer-Aided Design (ICCAD)*. IEEE, 2011.
- 466 [45] D. Ma, R. Thapa, X. Wang, X. Jiao, and C. Hao, “Workload-aware approximate computing
467 configuration,” in *2021 Design, Automation & Test in Europe Conference & Exhibition (DATE)*.
468 IEEE, 2021.
- 469 [46] C. Li, W. Luo, S. S. Sapatnekar, and J. Hu, “Joint precision optimization and high level synthesis
470 for approximate computing,” in *Proceedings of the 52nd Annual Design Automation Conference*,
471 2015.
- 472 [47] Y. Bengio, N. Léonard, and A. Courville, “Estimating or propagating gradients through stochas-
473 tic neurons for conditional computation,” *arXiv preprint arXiv:1308.3432*, 2013.
- 474 [48] E. Jang, S. Gu, and B. Poole, “Categorical reparameterization with gumbel-softmax,” *Internat-
475 ional Conference on Learning Representations*, 2017.
- 476 [49] C. J. Maddison, A. Mnih, and Y. W. Teh, “The concrete distribution: A continuous relaxation of
477 discrete random variables,” *International Conference on Learning Representations*, 2017.
- 478 [50] Y.-D. Kwon, J. Choo, B. Kim, I. Yoon, Y. Gwon, and S. Min, “Pomo: Policy optimization
479 with multiple optima for reinforcement learning,” *Advances in Neural Information Processing*
480 *Systems*, vol. 33, 2020.
- 481 [51] R. Wang, Z. Hua, G. Liu, J. Zhang, J. Yan, F. Qi, S. Yang, J. Zhou, and X. Yang, “A bi-level
482 framework for learning to solve combinatorial optimization on graphs,” *Advances in Neural*
483 *Information Processing Systems*, vol. 34, 2021.
- 484 [52] Y. Nandwani, D. Jindal, P. Singla *et al.*, “Neural learning of one-of-many solutions for com-
485 binatorial problems in structured output spaces,” in *International Conference on Learning*
486 *Representations*, 2021.
- 487 [53] A. Kumar and S. Levine, “Model inversion networks for model-based optimization,” *Advances*
488 *in Neural Information Processing Systems*, vol. 33, 2020.
- 489 [54] D. Brookes, H. Park, and J. Listgarten, “Conditioning by adaptive sampling for robust design,”
490 in *International Conference on Machine Learning*. PMLR, 2019.
- 491 [55] B. Trabucco, A. Kumar, X. Geng, and S. Levine, “Conservative objective models for effective
492 offline model-based optimization,” in *International Conference on Machine Learning*. PMLR,
493 2021.
- 494 [56] A. Kumar, A. Yazdanbakhsh, M. Hashemi, K. Swersky, and S. Levine, “Data-driven offline
495 optimization for architecting hardware accelerators,” 2022.
- 496 [57] V. Konda and J. Tsitsiklis, “Actor-critic algorithms,” *Advances in Neural Information Processing*
497 *Systems*, vol. 12, 1999.
- 498 [58] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu,
499 “Asynchronous methods for deep reinforcement learning,” in *International Conference on*
500 *Machine Learning*. PMLR, 2016.

- 501 [59] M. Paulus, D. Choi, D. Tarlow, A. Krause, and C. J. Maddison, “Gradient estimation with
502 stochastic softmax tricks,” *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- 503 [60] K. Struminsky, A. Gadetsky, D. Rakitin, D. Karpushkin, and D. P. Vetrov, “Leveraging recursive
504 gumbel-max trick for approximate inference in combinatorial spaces,” *Advances in Neural
505 Information Processing Systems*, vol. 34, 2021.
- 506 [61] M. Fey and J. E. Lenssen, “Fast graph representation learning with PyTorch Geometric,” in
507 *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- 508 [62] M. Wang, D. Zheng, Z. Ye, Q. Gan, M. Li, X. Song, J. Zhou, C. Ma, L. Yu, Y. Gai, T. Xiao,
509 T. He, G. Karypis, J. Li, and Z. Zhang, “Deep graph library: A graph-centric, highly-performant
510 package for graph neural networks,” *arXiv preprint arXiv:1909.01315*, 2019.
- 511 [63] S. Mirjalili, “Genetic algorithm,” in *Evolutionary algorithms and neural networks*. Springer,
512 2019.
- 513 [64] D. Whitley, “A genetic algorithm tutorial,” *Statistics and computing*, vol. 4, no. 2, 1994.
- 514 [65] M. V. Pogančić, A. Paulus, V. Musil, G. Martius, and M. Rolínek, “Differentiation of blackbox
515 combinatorial solvers,” in *International Conference on Learning Representations*, 2019.
- 516 [66] L. Deng, “The mnist database of handwritten digit images for machine learning research,” *IEEE
517 Signal Processing Magazine*, vol. 29, no. 6, 2012.
- 518 [67] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in
519 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- 520 [68] W. Hamilton, Z. Ying, and J. Leskovec, “Inductive representation learning on large graphs,”
521 *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- 522 [69] AMD/Xilinx, “Vitis ai development environment,” [https://www.xilinx.com/products/
523 design-tools/vitis/vitis-ai.html](https://www.xilinx.com/products/design-tools/vitis/vitis-ai.html).
- 524 [70] —, “Vivado development tool,” <https://www.xilinx.com/products/design-tools/vivado.html>.
- 525 [71] P. J. Huber, “Robust estimation of a location parameter,” in *Breakthroughs in statistics*. Springer,
526 1992, pp. 492–518.
- 527 [72] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin,
528 N. Gimeshain, L. Antiga *et al.*, “Pytorch: An imperative style, high-performance deep learning
529 library,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- 530 [73] M. Fey and J. E. Lenssen, “Fast graph representation learning with pytorch geometric,” *arXiv
531 preprint arXiv:1903.02428*, 2019.
- 532 [74] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *International
533 Conference on Learning Representations*, 2015.
- 534 [75] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, “Neural message passing
535 for quantum chemistry,” in *International Conference on Machine Learning*. PMLR, 2017, pp.
536 1263–1272.
- 537 [76] G. Corso, L. Cavalleri, D. Beaini, P. Liò, and P. Veličković, “Principal neighbourhood aggrega-
538 tion for graph nets,” *Advances in Neural Information Processing Systems*, vol. 33, 2020.

539 **Checklist**

- 540 1. For all authors...
- 541 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s
542 contributions and scope? [Yes]
- 543 (b) Did you describe the limitations of your work? [Yes] See Sec. 5.
- 544 (c) Did you discuss any potential negative societal impacts of your work? [Yes] See
545 Appendix D.
- 546 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
547 them? [Yes]
- 548 2. If you are including theoretical results...
- 549 (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Sec 2 and
550 Sec 3
- 551 (b) Did you include complete proofs of all theoretical results? [Yes] See Appendix A.1,
552 A.2 and A.3
- 553 3. If you ran experiments...
- 554 (a) Did you include the code, data, and instructions needed to reproduce the main experi-
555 mental results (either in the supplemental material or as a URL)? [Yes]
- 556 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
557 were chosen)? [Yes] See the experiment details in Appendix C.
- 558 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
559 ments multiple times)? [Yes]
- 560 (d) Did you include the total amount of compute and the type of resources used (e.g., type
561 of GPUs, internal cluster, or cloud provider)? [Yes] See Appendix C.
- 562 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 563 (a) If your work uses existing assets, did you cite the creators? [Yes] The dataset in
564 application II is adopted from [30].
- 565 (b) Did you mention the license of the assets? [Yes] See Appendix E.
- 566 (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
567 All the new assets are publicly available.
- 568 (d) Did you discuss whether and how consent was obtained from people whose data you’re
569 using/curating? [Yes] All the assets obtained from others are publicly available.
- 570 (e) Did you discuss whether the data you are using/curating contains personally identifiable
571 information or offensive content? [Yes] See Appendix E and our data contains no
572 human information or offensive content.
- 573 5. If you used crowdsourcing or conducted research with human subjects...
- 574 (a) Did you include the full text of instructions given to participants and screenshots, if
575 applicable? [N/A]
- 576 (b) Did you describe any potential participant risks, with links to Institutional Review
577 Board (IRB) approvals, if applicable? [N/A]
- 578 (c) Did you include the estimated hourly wage paid to participants and the total amount
579 spent on participant compensation? [N/A]

580 A Deferred Theoretical Arguments

581 A.1 Proof of Theorem 1

582 We analyze the rounding process of the relaxed solution $\bar{X} = \mathcal{A}_\theta(C)$, $\bar{X} \in [0, 1]^n$ into the integral
 583 solution $\hat{X} \in \{0, 1\}^n$. Let $\bar{X}_i, \hat{X}_i, i = \{1, 2, \dots, n\}$ denote their entries respectively. The rounding
 584 procedure has no requirement on the order of the rounding sequence, w.l.o.g, suppose we round from
 585 index $i = 1$ to $i = n$. In practice, it might be better to sort \bar{X} and round the entries according to their
 586 ranks. We have the following inequations:

$$\begin{aligned}
 & l_r(\theta; C) \\
 &= f_r([\bar{X}_1, \bar{X}_2, \dots, \bar{X}_n]; C) + \beta g_r([\bar{X}_1, \bar{X}_2, \dots, \bar{X}_n]; C) \\
 &\stackrel{(a)}{\geq} \bar{X}_1 (f_r([1, \bar{X}_2, \dots, \bar{X}_n]; C) + \beta g_r([1, \bar{X}_2, \dots, \bar{X}_n]; C)) \\
 &\quad + (1 - \bar{X}_1) (f_r([0, \bar{X}_2, \dots, \bar{X}_n]; C) + \beta g_r([0, \bar{X}_2, \dots, \bar{X}_n]; C)) \\
 &\geq \bar{X}_1 \left(\min_{j_1 \in \{0,1\}} f_r([j_1, \bar{X}_2, \dots, \bar{X}_n]; C) + \beta g_r([j_1, \bar{X}_2, \dots, \bar{X}_n]; C) \right) \\
 &\quad + (1 - \bar{X}_1) \left(\min_{j_1 \in \{0,1\}} f_r([j_1, \bar{X}_2, \dots, \bar{X}_n]; C) + \beta g_r([j_1, \bar{X}_2, \dots, \bar{X}_n]; C) \right) \\
 &\stackrel{(b)}{=} (f_r([\hat{X}_1, \bar{X}_2, \dots, \bar{X}_n]; C) + \beta g_r([\hat{X}_1, \bar{X}_2, \dots, \bar{X}_n]; C)) \tag{9} \\
 &\geq \min_{j_2 \in \{0,1\}} (f_r([\hat{X}_1, j_2, \dots, \bar{X}_n]; C) + \beta g_r([\hat{X}_1, j_2, \dots, \bar{X}_n]; C)) \\
 &= (f_r([\hat{X}_1, \hat{X}_2, \dots, \bar{X}_n]; C) + \beta g_r([\hat{X}_1, \hat{X}_2, \dots, \bar{X}_n]; C)) \\
 &\geq \dots \\
 &\geq \min_{j_n \in \{0,1\}} (f_r([\hat{X}_1, \hat{X}_2, \dots, j_n]; C) + \beta g_r([\hat{X}_1, \hat{X}_2, \dots, j_n]; C)) \\
 &= f_r(\hat{X}; C) + \beta g_r(\hat{X}; C) \\
 &\stackrel{(c)}{=} f(\hat{X}; C) + \beta g(\hat{X}; C),
 \end{aligned}$$

587 where (a) is due to $l_r(\theta; C)$'s entry-wise concavity w.r.t. \bar{X} and Jensen's inequality, (b) is due
 588 to the definition $\hat{X}_1 = \arg \min_{j_1 \in \{0,1\}} f_r([j_1, \bar{X}_2, \dots, \bar{X}_n]; C) + \beta g_r([j_1, \bar{X}_2, \dots, \bar{X}_n]; C)$, and (c)
 589 is due to the assumption that the neural network based proxies could learn the objective and the
 590 constraints perfectly for $\hat{X} \in \{0, 1\}^n$. The inequalities above demonstrate the fact that the loss value
 591 is non-increasing via the whole rounding process. By this, once the learnt parameter θ achieves
 592 $l_r(\theta; C) < \beta$, we could get $f(\hat{X}; C) + \beta g(\hat{X}; C) \leq l_r(\theta; C) < \beta$. Because of the settings that
 593 $f(\cdot), g(\cdot) \geq 0$, we have $f(\hat{X}; C) < l_r(\bar{X}; C)$, s.t. $g(\hat{X}; C) < 1$.

594 A.2 Proof of Theorem 2

595 Set $h_r(\bar{X}) = \sum_{X \in \{0,1\}^n} h(X) \prod_{j=1}^n \bar{X}_j^{X_j} (1 - \bar{X}_j)^{(1-X_j)}$.

596 We first prove that $h_r(\bar{X})$ with the form above satisfies (a) $h_r(X) = h(X)$ for $X \in \{0, 1\}^n$.

597 Given one $X' \in \{0, 1\}^n$, by setting $\bar{X} = X'$, we have

$$\prod_{j=1}^n X_j'^{X_j'} (1 - X_j')^{(1-X_j')} = 1, \quad \text{if } X = X', \text{ and otherwise } 0.$$

598 Therefore, in $h_r(X')$, there is only one term $h(X') \prod_{j=1}^n X_j'^{X_j'} (1 - X_j')^{(1-X_j')} = h(X')$ left. So,
 599 $h_r(X') = h(X')$, which satisfies (a).

600 Then we prove that $h_r(\bar{X})$ satisfies (b) h_r is entry-wise affine. From the definition, we have:

$$\bar{X}_j^{X_j} (1 - \bar{X}_j)^{1-X_j} = \begin{cases} \bar{X}_j & X_j = 1 \\ 1 - \bar{X}_j & X_j = 0. \end{cases}$$

601 Consider two sequences $\bar{X}, \bar{X}' \in \{0, 1\}^n$ with the entries $\bar{X}_i, \bar{X}'_i, i = \{1, 2, \dots, n\}$. We have

$$\begin{aligned} & \gamma[\bar{X}_i^{X_i}(1 - \bar{X}_i)^{1-X_i}] + (1 - \gamma)[\bar{X}'_i^{X_i}(1 - \bar{X}'_i)^{1-X_i}] \\ &= \begin{cases} 1 - \gamma\bar{X}_i - (1 - \gamma)\bar{X}'_i & X_i = 0 \\ \gamma\bar{X}_i + (1 - \gamma)\bar{X}'_i & X_i = 1. \end{cases} \\ &= [\gamma\bar{X}_i + (1 - \gamma)\bar{X}'_i]^{X_i} [1 - \gamma\bar{X}_i - (1 - \gamma)\bar{X}'_i]^{(1-X_i)}. \end{aligned}$$

602 W.l.o.g, we assume that only the entries \bar{X}_i and \bar{X}'_i in the two sequences are different. For any
603 $\gamma \in [0, 1]$, we may use the above equality and have

$$\begin{aligned} & \gamma h_r(\bar{X}) + (1 - \gamma)h_r(\bar{X}') \\ &= \gamma \sum_{X \in \{0,1\}^n} h(X) \prod_{j=1}^n \bar{X}_j^{X_j} (1 - \bar{X}_j)^{(1-X_j)} + (1 - \gamma) \sum_{X \in \{0,1\}^n} h(X) \prod_{j=1}^n \bar{X}'_j^{X_j} (1 - \bar{X}'_j)^{(1-X_j)} \\ &= \gamma \sum_{X \in \{0,1\}^n} h(X) \bar{X}_i^{X_i} (1 - \bar{X}_i)^{1-X_i} \prod_{j=1, j \neq i}^n \bar{X}_j^{X_j} (1 - \bar{X}_j)^{(1-X_j)} \\ &+ (1 - \gamma) \sum_{X \in \{0,1\}^n} h(X) \bar{X}'_i^{X_i} (1 - \bar{X}'_i)^{1-X_i} \prod_{j=1, j \neq i}^n \bar{X}_j^{X_j} (1 - \bar{X}_j)^{(1-X_j)} \\ &= \sum_{X \in \{0,1\}^n} h(X) \left[\gamma \bar{X}_i^{X_i} (1 - \bar{X}_i)^{1-X_i} + (1 - \gamma) \bar{X}'_i^{X_i} (1 - \bar{X}'_i)^{1-X_i} \right] \prod_{j=1, j \neq i}^n \bar{X}_j^{X_j} (1 - \bar{X}_j)^{(1-X_j)} \\ &= \sum_{X \in \{0,1\}^n} h(X) [\gamma\bar{X}_i + (1 - \gamma)\bar{X}'_i]^{X_i} [1 - \gamma\bar{X}_i - (1 - \gamma)\bar{X}'_i]^{(1-X_i)} \prod_{j=1, j \neq i}^n \bar{X}_j^{X_j} (1 - \bar{X}_j)^{(1-X_j)} \\ &= h_r(\gamma\bar{X} + (1 - \gamma)\bar{X}') \end{aligned}$$

604 Thus, we prove that the form of $h_r(X)$ is entry-wise affine.

605 A.3 Proof of Proposition 1

606 Suppose that the output of function $h(X_1, X_2)$ is denoted as follows.

$$a_0 = h(0, 0), \quad a_1 = h(0, 1), \quad a_2 = h(1, 0), \quad a_3 = h(1, 1) \quad (10)$$

607 Then, we pick out the largest value a_i among a_0, a_1, a_2, a_3 . W.l.o.g, we assume that a_0 is the largest
608 and they hold the following inequations:

$$a_0 \geq a_1, \quad a_0 \geq a_2, \quad a_0 \geq a_3. \quad (11)$$

609 Then, we define our entry-wise concave function $h_r(X_1, X_2)$ as follows.

$$\begin{aligned} h_r(X_1, X_2) &= a_0 - \sum_{i=1}^3 \text{Relu}(f[i]) \\ &= a_0 - \text{Relu}(f[1]) - \text{Relu}(f[2]) - \text{Relu}(f[3]), \end{aligned} \quad (12)$$

610 where

$$\begin{aligned} \text{Relu}(f[1]) &= \text{Relu}((a_0 - a_1)(X_2 - X_1)) = \begin{cases} a_0 - a_1 & X_1 = 0, X_2 = 1 \\ 0 & \text{otherwise,} \end{cases} \\ \text{Relu}(f[2]) &= \text{Relu}((a_0 - a_2)(X_1 - X_2)) = \begin{cases} a_0 - a_2 & X_1 = 1, X_2 = 0 \\ 0 & \text{otherwise,} \end{cases} \\ \text{Relu}(f[3]) &= \text{Relu}((a_0 - a_3)(X_1 + X_2 - 1)) = \begin{cases} a_0 - a_3 & X_1 = 1, X_2 = 1 \\ 0 & \text{otherwise.} \end{cases} \end{aligned} \quad (13)$$

611 B Additional Results

612 B.1 Application I

613 Here, we display some additional visualization results of the feature-based edge covering and node
 614 matching problems in application I in Fig. 5. In both edge covering and node matching problems,
 615 our method based on entry-wise affine proxies could avoid the multiplication between some large
 616 numbers so that the final cost could be low enough. For example, our method avoids $85*85$ in the
 617 first row, $97*76$ in the second row, $71*91$ in the third row, and $69*98$ in the fourth row, which are all
 selected by the method with Gumbel-Softmax tricks.

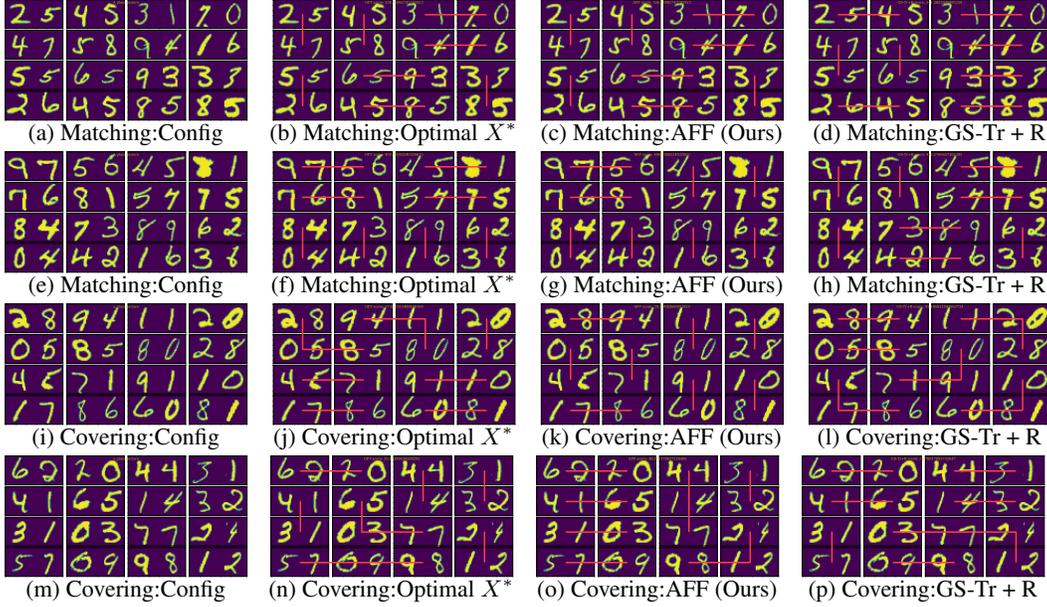


Figure 5: The additional visualization of the edge covering and node matching for Application I.

618

619 B.2 Application II

620 The additional visualization results of DSP-LUT usage amount relationship on the test configurations is
 621 shown in Figure. 6. Our entry-concave proxies generates the lowest LUT-DSP combinations among all
 622 the methods. To be fair, we also pick the best results from 200 randomly sampled HLS tool’s assignment
 623 as a baseline (called HLS), to show the superiority of the other optimization methods. The GS-Tr+R and
 624 the RL method outperforms the HLS baseline, while the generic methods SA and GA only marginally out-
 625 perform and are comparable with the best of HLS random solutions.
 626
 627
 628
 629
 630
 631

632 The average ranking of the LUT usage under the constraint of the maximum DSP usage amount
 633 is shown in Table. 6, which adds to another two baselines, the ‘Naive’ baseline and the GS-Tr+S
 634 method. It turns out that these two methods could not generate feasible results that satisfy the DSP
 635 usage threshold when the threshold is relatively low, thus we put them in the last place in the ranking
 636 if they could not generate feasible solutions.

637 We also include the percentages of cases where each method takes the first, second and third places
 638 according to the rank, which is shown in Table 5. The comparison about how different proxies

Method	1-st (%)	2-nd (%)	3-rd (%)
SA	12.3	10.3	13.4
GA	14.8	15.7	18.2
RL	18.5	21.8	19.5
GS-Tr+R	20.9	27.2	22.7
CON	39.5	24.8	26.0

Table 5: The percentage that each method occupies in the first, second and third place of the ranking.

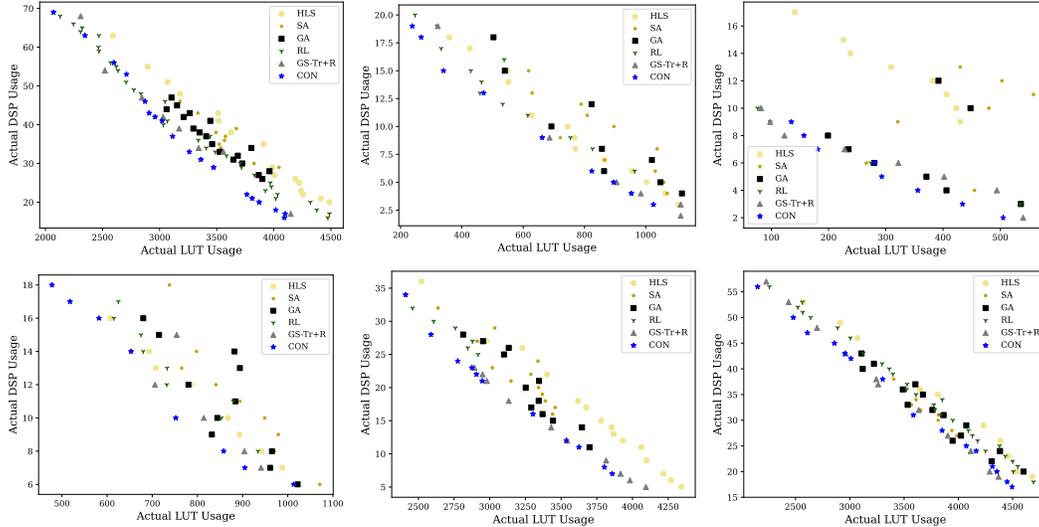


Figure 6: The additional visualization of LUT-DSP usage relationship. The HLS baseline denotes the optimal synthesis results among 200 random mappings.

DSP usage	40%	45%	50%	55%	60%	65%	70%	rank-avg
SA	3.94	3.68	3.99	3.96	4.50	5.03	5.21	4.33
GA	2.88	3.09	3.73	3.56	4.20	4.03	4.42	3.70
Naïve	6.88	6.62	6.78	6.85	6.76	6.75	5.90	6.64
RL	3.50	4.19	4.33	3.73	3.22	3.14	2.64	3.53
GS-Tr+S	5.02	4.97	3.83	3.90	3.37	3.35	3.67	4.01
GS-Tr+R	3.93	3.39	2.94	3.57	2.52	2.98	3.23	3.22
CON	2.15	2.26	2.46	2.47	3.43	2.72	2.93	2.63

Table 6: The result of LUT DSP balancing problem in application II. The DSP usage thresholds are from 40% to 70%. For the Naïve and GS-Tr+S method, if there are no feasible results under the DSP usage constraint, we put them in the last place.

639 approximate the ground-truth LUT usage amount is shown in Fig 7. Again, the entry-wise affine
 640 proxy may introduce large error while the entry-wise concave proxy could approximate in a better
 641 sense.

642 We investigate the training time of RL method, GS-Tr method and the entry-wise concave method.
 643 The comparison among these methods is shown in Fig 8. We run all the three methods on the same
 644 server with 2 Intel(R) Xeon(R) Gold 6248R CPUs, 1000GB RAM in total. In each experiment we
 645 take 26 processes of the CPU and run on one Quadro RTX 6000 GPU card. We count the time cost
 646 during training and select the models at different epochs for testing at intervals. Note that the training
 647 objectives in all methods use proxies while the testing results are the outputs given by the HLS tool.
 648 Due to the fact that inferring via HLS consumes a lot of time, we only test limited numbers of cases
 649 to draw the figure and thus the curves are not smooth.

650 B.3 Application III

651 For each method, we count the ratio that the relative error based on the assignment given by this
 652 method exceeds the optimal assignment (ratio = relative error / OPT relative error - 1). The smaller
 653 the ratio is, the closer the method’s relative error is to the optimal solution. The results are shown
 654 in Table 7. We also include the training time that the methods require to achieve the corresponding
 655 performance in the table. All the methods run on the same server with 2 Intel(R) Xeon(R) Gold
 656 6248R CPUs, 1000GB RAM in total. In each experiment we take 26 processes of the CPU and run
 657 on one Quadro RTX 6000 GPU card. The time is obtained by counting the least epoch that a model

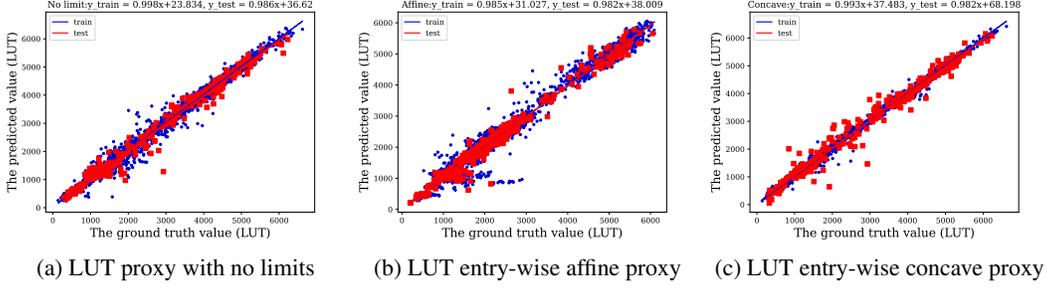


Figure 7: The visualization of different proxies in LUT learning in application II.

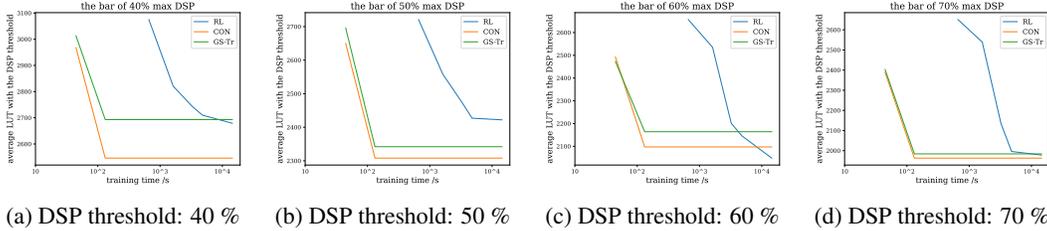


Figure 8: Training time of RL, vs GS-Tr, vs CON with different constraints on DSP usage.

658 achieves its reported performance, the time of OPT is the time that the brute-force search takes in the
 659 testset.

Threshold θ	C-In	C-Out	Naive	RL	GS-Tr+S	GS-Tr+R	CON	AFF	OPT
3	348.47	349.09	30.68	282.31	75.68	16.78	14.80	11.91	0
5	209.70	209.07	30.80	159.07	69.17	23.53	8.2	13.50	0
8	99.41	99.06	29.90	77.21	47.68	24.03	18.80	17.28	0
Time / s	0	0	90+	9851+	90+	90+	92+	92+	32776+

Table 7: The averaged ratios (%) that the relative errors of different methods exceed the OPT on application III. The required training time to achieve the performance is listed at the bottom row.

660 B.4 Study I: Further Evaluation on The Learning Capability of Different Proxies.

661 We have studied the learning capability of different proxies for Application II, as shown in Fig. 7 and Fig. 3. We
 662 further study how proxies under different constraints fit with historical data on application III. We show the Huber
 663 loss [71] of the proxies in Table 8. The objective of application III has 4-order moment, the three proxies obtain
 664 comparable approximation performance.
 665
 666
 667

Proxy	AFF	CON	W/O limit
Train	0.19	0.17	0.10
Test	0.19	0.18	0.19

Table 8: Huber loss on application III.

668 B.5 Study II: The Effectiveness of Our ‘Rounding’ Process

669 Here we empirically study the effectiveness of our rounding procedure. We show the average
 670 relaxed loss value ($l_r = f_r(\bar{X}; C) + \beta g_r(\bar{X}; C)$, $\bar{X} \in [0, 1]^n$), the average rounded loss value
 671 ($l'_r = f_r(\hat{X}; C) + \beta g_r(\hat{X}; C)$, $\hat{X} \in \{0, 1\}^n$) and the average true value of the rounded assignment
 672 ($l = f(\hat{X}; C) + \beta g(\hat{X}; C)$) on the testset of each problem in Table 9. According to the table, we
 673 observe that both the methods that adopt entry-wise affine proxies and entry-wise concave proxies are
 674 guaranteed to obtain a drop of the loss values after our rounding procedure. However, for the proxies
 675 that do not satisfy the constraints, the Naïve method and the GS-Tr-R baseline could not always
 676 guarantee such a drop after the rounding process. In particular, the rounding in GS-Tr-R increases the

	Edge covering (App. I)				Node matching (App. I)				Resource allocation (App. II)			A×C circuit design (App. III)			
Proxy	Naïve	GS-Tr-R	CON	AFF	Naïve	GS-Tr-R	CON	AFF	Naïve	GS-Tr-R	CON	Naïve	GS-Tr-R	CON	AFF
l_r	62.58	78.56	80.91	46.37	5316.07	13103.39	5922.61	5389.95	5899.10	3485.60	2785.95	4.41	4.23	6.75	6.67
l'_r	70.20	51.04	52.09	45.73	442.15	442.03	430.11	432.63	3350.35	2741.32	2599.89	9.31	7.27	5.22	5.19
l	68.52	46.91	52.04	44.55	429.12	429.39	422.47	418.96	2901.19	2749.08	2511.70	6.98	6.57	6.21	6.17
OPT		42.69				416.05				no OPT				5.36	

Table 9: The relaxed loss value l_r , the rounded loss value l'_r and its true value l of the methods.

677 loss in the application of A×C circuit design, while the rounding in Naïve increases the loss in the
678 applications of edge covering and A×C circuit design.

679 C Experimental details

680 All of the experiments are carried out on the same server with 2 Intel(R) Xeon(R) Gold 6248R
681 CPUs, 1000GB RAM in total. In each experiment we take 26 processes of the CPU and run on one
682 Quadro RTX 6000 GPU card. The maximum GRAM of the Quadro RTX 6000 GPU is 24GB. The
683 proxies that satisfies our principle (AFF, CON) and GS-Tr run on PyTorch [72] frame with PyTorch
684 geometric [73]. The RL baseline follow the actor-critic technique in [57]. [30] also utilizes the
685 same RL technique to solve the same problem. The details of each dataset is displayed in Table. 10.
686 Adam [74] is used as the optimizer in all of the experiments. All the experiment results are conduct
687 and averaged under three random seeds 12345, 23456 and 34567.

688 To be fair, in the training process, we first train the baseline methods, such as the proxy without
689 constraints and \mathcal{A}_θ based on the Gumbel-softmax trick. Then we train our entry-wise concave proxy
690 and \mathcal{A}_θ with exactly the same hyper-parameters with the baselines except for the necessary changes
691 to construct the entry-wise concavity.

Task	Toy example	Application I	Application II	Application III
f_r, g_r training	95,000	95,000	7,200	95,000
f_r, g_r testing	5,000	5,000	800	5,000
A_θ training	10,000	10,000	40	1,000
A_θ testing	500	500	20	500

Table 10: The number of instances in each dataset.

692 C.1 The toy example

693 The ground truth of the objectives is designed as follows.

$$f(X_1, X_2; C) = g_1(C)X_1 + g_2(C)X_2 + g_3(C)X_1X_2 + g_4(C), \quad (14)$$

694 where $C = [C_1, C_2]$ and

$$\begin{aligned} g_1 &= (580 - 10C_1 - 3C_2)/33, \\ g_2 &= (580 - 10C_2 - 3C_1)/33, \\ g_3 &= (3C_1 + 3C_2)/45, \\ g_4 &= -(5C_1 + 5C_2)/33 + 60. \end{aligned} \quad (15)$$

695 The constants are set arbitrarily. To match our graph-based pipeline, in this toy example, we also
696 build a single edge graph for each configuration where the two nodes are associated with the attributes
697 C_1 and C_2 , and the binary variables X_1 and X_2 .

698 **The proxy without constraints:** We use 3 layers of GraphSAGE [68] convolutional layers that take
699 both the node attributes C and the optimization variables \bar{X} as inputs with leaky ReLU activation and
700 batch normalization. Then, the structure is followed by a global mean pooling. After several MLP
701 layers, the proxy outputs the cost. We use MSE loss to train this proxy. The learning rate is set as
702 $1e-2$. The batch size is set as 4096. The reported performance is trained within 200 epochs.

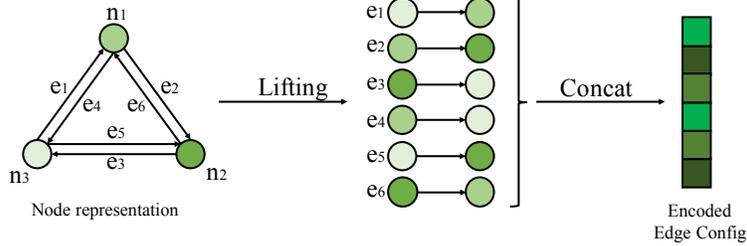


Figure 9: The generalization from node problems to edge problems.

703 **The entry-wise concave proxy:** We use 3 layers of GraphSAGE with the same hyper-parameters
 704 as the proxy without constraints but it only takes the configuration C as inputs to encode the
 705 configuration. Then, the encoded configuration $H \in \mathbb{R}^{|V| \times d}$ (d is the dimension of the latent
 706 node feature) is separated equally into two parts $U \in \mathbb{R}^{|V| \times \frac{d}{2}}$ and $W \in \mathbb{R}^{|V| \times \frac{d}{2}}$, we construct
 707 $U\bar{X} + W$, ($\bar{X} \in [0, 1]^{|V|}$), then we multiply $U_1\bar{X}_1 + W_1$ with $U_2\bar{X}_2 + W_2$ to obtain the latent node
 708 representation $\phi(\bar{X}; C)$, followed by a linear layer as the implementation of the AFF proxy $h_r^a(\bar{X}; C)$
 709 as introduced in Section 4 Eq. 6. MSE loss is used as the criterion, the learning rate is 1e-2, the batch
 710 size is 4096, the model is trained within 200 epochs.

711 \mathcal{A}_θ **based on Gumbel-softmax tricks:** When training \mathcal{A}_θ , we also use 3 GraphSAGE layers to
 712 encode the configuration C with leaky Relu activation and batch normalization. Then the encoded
 713 latent feature is followed by fully connected layers to reduce the dimension and the Gumbel-softmax
 714 trick to sample a distribution from the soft probability predicted by the model. We use the soft
 715 Gumbel-softmax [48, 49] without the straight through trick. The learning rate is set as 1e-2, the batch
 716 size is 4096, the model is trained for 200 epochs.

717 \mathcal{A}_θ **in {relaxation (Naïve), relaxation with entry-wise concave proxy (Ours)}:** In \mathcal{A}_θ , the structure
 718 is the same as that based on Gumbel-softmax tricks except that the structure has no pooling layer and
 719 takes a Sigmoid layer. The learning rate is 1e-2, the batch size is 4096, the model is trained for 200
 720 epochs.

721 C.2 Application I - Edge Covering

722 **Dataset details:** Each configuration C is a 4×4 grid graph. Each node of the graph consists of two
 723 images randomly selected from the MNIST dataset [66] and thus represents a number between 00
 724 and 99.

725 The ground truth of the objective are designed as follows:

$$f(X; C) = \sum_{e \in E} w_e X_e, \quad (16)$$

726 where

$$w_e = (C_v + C_u)/3 + (C_v C_u)/100, \text{ for } e = (u, v) \in E. \quad (17)$$

727 **The proxy without constraints:** We firstly utilize ResNet-50 [67] to extract the latent fixed node
 728 feature and then send the feature into a GNN, the GNN is also based on 3 MPNN [75] layers, which
 729 involves the edge assignment \bar{X} , ($\bar{X} \in [0, 1]^{|E|}$) in the message passing. Global mean pooling is
 730 used to generate the final predicted value $f_r(\bar{X}; C)$. MSE loss is utilized as the criterion, the learning
 731 rate is 5e-3, the batch size is 160.

732 **The entry-wise affine proxy:** We also use ResNet-50 and 3 layers of MPNN [75] which take the
 733 output of ResNet-50 as inputs with leaky Relu activation and batch normalization to encode the
 734 configuration C into the latent node representation $H' \in \mathbb{R}^{|V| \times d_1}$ (d_1 is the dimension of the node
 735 feature). After the encoding procedure, the encoded node features are lifted to each side of the edges
 736 according to the edge index, then these node features on two sides of the edges are concatenated
 737 together and sent into MLP layers to generate the latent edge representation $H'' \in \mathbb{R}^{|E| \times d_2}$ (d_2

738 is the dimension of the edge feature). Then H'' is separated into two parts $U \in \mathbb{R}^{|E| \times \frac{3d_2}{4}}$ and
739 $W \in \mathbb{R}^{|E| \times \frac{d_2}{4}}$, and we calculate $U\bar{X} + W$, ($\bar{X} \in [0, 1]^{|E|}$) to construct the latent representation
740 $\phi(\bar{X}; C)$. The structure is followed by mean pooling and linear layers to construct the AFF proxy
741 $h_r^a(\bar{X}; C)$. The whole procedure generalizes our framework from solving node problems to edge
742 problems, as is shown in Fig. 9. We use MSE loss for training, the learning rate is set as 5e-3, the
743 batch size is 160.

744 **The entry-wise concave proxy:** The network shares the same structure with the AFF proxy in
745 the front part, while we utilize linear layers mixed with a negative Relu function to construct the
746 CON proxy $h_r^c(\bar{X}; C) = \langle w^c, -\text{Relu}(\phi(\bar{X}; C)) \rangle + b$, as introduced in Section 4. Note that we use
747 torch.clamp() function to control the entries in w^c greater or equal to zero in each batch of data during
748 the training process. We use MSE loss for training and set the learning rate as 5e-3, the batch size is
749 160.

750 **The RL baseline:** We apply an actor critic model [57]. This model consists of 4 key components:
751 1)States, the states are formulated as every possible partially assigned grid graph; 2) Actions, given
752 the current state and the currently candidate edges of the grid graph, the action is which new edge to
753 pick. Note that the model is only allowed to pick from the edges which connect at least one node
754 that has not been covered yet; 3) State transition, given a state and an action, the probability of the
755 next states; 4)Reward, the reward is 0 for all intermediate actions, in the last action the reward is the
756 evaluation of the covering score predicted by the proxy without constraints.

757 In each state at a time step, we extract the features from the last layer of the proxy without constraints
758 f_r . We utilize another ResNet-50 + GNN to encode the whole grid graph into a vector encoding.
759 The features are further combined with the vector embedding as the state encoding. Then the state
760 encoding is sent into the policy network that is made up of multiple MLP layers to output the critic
761 value c and the action a which indicates the next edge to pick from. The loss for the actor is calculated
762 by subtracting the reward by c , and we use Huber loss to make c close to the reward. In each state, the
763 model would only choose from the edges which connect at least one node that has not been covered
764 yet. The reward is defined as the negative proxy prediction:

$$r_t = \begin{cases} -f_r(X; C), & s = T \\ 0, & 0 < s < T, \end{cases} \quad (18)$$

765 where T is the max step, and s denotes the number of the step. The learning rate is set as 1e-2, the
766 discount factor for the reward is set as 0.95, we train the RL baseline for more than 20,000 epochs to
767 achieve the reported performance.

768 **The constraint $g_r(\bar{X}; C)$:** As to the penalty constraint $g_r(\bar{X}; C) = \sum_{v \in V} \prod_{e: v \in e} (1 - \bar{X}_e)$ which
769 naturally satisfies our definition of CO problems in Section 2, we apply the log-sum-exponential trick
770 $\sum_{v \in V} \exp(\sum_{v \in e} \log(1 - \bar{X}_e))$ to calculate it via message passing in PyTorch geometric.

771 **\mathcal{A}_θ based on Gumbel-softmax trick:** We utilize 3 GraphSAGE layers to encode the node feature,
772 with leaky Relu activation function and batch normalization. The encoded node features are lifted to
773 each side of the edges, concatenated together and then sent into MLP layers to reduce the dimension
774 and map to $\mathcal{A}_\theta(C) \in [0, 1]^{|E|}$. Then the model is followed by the Gumbel-softmax trick to obtain the
775 output $X \sim \text{Ber}(\mathcal{A}_\theta)$. We use the soft Gumbel-softmax [48, 49] without the straight through trick.
776 The learning rate is set as 1e-3, the batch size is 60.

777 **\mathcal{A}_θ in {relaxation (Naïve), relaxation with entry-wise concave proxy (Ours)}:** The model shares
778 the same structure as that based on the Gumbel-softmax trick, except that the Gumbel-softmax trick
779 is replaced by $\bar{X} \in [0, 1]^{|E|}$ directly. The learning rate is set as 1e-3, the batch size is 60.

780 C.3 Application I - Node Matching

781 The ground truth of the objective are designed as follows:

$$f(X; C) = \sum_{e \in E} w_e X_e, \quad (19)$$

782 where

$$w_e = C_v C_u, \text{ for } e = (u, v) \in E. \quad (20)$$

783 Every structure design keeps the same as Application I - Edge Covering except for the extra penalty
 784 constraint $\prod_{\substack{e_1, e_2: v \in e_1, e_2 \\ e_1 \neq e_2}} \bar{X}_{e_1} \bar{X}_{e_2}$ in Eq. (8) which naturally follows our request on g_r , and is also
 785 conducted via matrix operations on PyTorch geometric. As to the RL baseline, the structure design
 786 is basically the same as Application I, while in this problem, at each time step, the model is only
 787 allowed to pick an edge whose two nodes are both not covered. The reward is defined as follows:

$$r_t = \begin{cases} -f_r(X; C) & s = T \\ -\beta & 0 < s < T, \text{ no options} \\ 0, & 0 < s < T, \text{ option exists,} \end{cases} \quad (21)$$

788 where “no options” means that there are some covered nodes whose neighboring nodes have been all
 789 covers, “option exists” denotes the case when eligible edges still exist, T is the max step, β is a large
 790 hyper-parameter, and s denotes the number of the step. The learning rate is set as 1e-2, the discount
 791 factor for the reward is set as 0.95, we train the RL baseline for more than 20,000 epochs to achieve
 792 the reported performance.

793 C.4 Application II - Resource Binding Optimization

794 **Dataset details** In this application, we focus on the resource binding problems in field-programmable
 795 gate array (FPGA) design. Each configuration C in the dataset is a data flow graph (DFG) with more
 796 than 100 nodes. Each node represents an arithmetic operation such as multiplication or addition. The
 797 operations need to be one-to-one mapped into a micro circuit to carry out the calculation. Given an
 798 assignment of the mapping, we run high-level synthesis (HLS) simulation tools to obtain the actual
 799 circuit resource usage under the assignment, which might take up to hours time. Note that different
 800 assignments of the mapping could result in vastly different actual resource usage.

801 In this dataset, we focus on the resource balancing problems between digital signal processors (DSP)
 802 and look-up tables (LUT). Here DSP is a small processor that is able to quickly perform mathematical
 803 operation on streaming digital signals, LUT is the small memory that is used to store truth tables
 804 and perform logic functions. The optimization goal of the dataset is to allocate those nodes with
 805 pragma to either LUT or DSP, such that the actual usage amount of LUT could be minimized given
 806 a maximum usage amount of the DSP usage. We use 1 (LUT) or 0 (DSP) to assign each node’s
 807 mapping. We encode the fixed node feature into a 10-dimension embedding which contains the
 808 following information: 1) 4 digits to indicate the types of nodes in {input, m-type, intermediate-type,
 809 output}; 2) 5 digit binary encoding of the node’s calculation precision, from 2 bits to 32 bits; 1
 810 digit encoding that indicates whether the node requires pragma. For those nodes that do not require
 811 pragma, HLS tools have a set of heuristic assignments to the nodes during the simulation.

812 **The proxy without constraints:** We separately utilize two GraphSAGE GNN models [68] to predict
 813 the LUT usage $f_r(\bar{X}; C)$ and the DSP usage $g_r(\bar{X}; C)$, the structure of them are the same. We
 814 use 3 layers of GraphSAGE [68] convolutional layers that take both the node attributes C and the
 815 optimization variables \bar{X} as inputs with leaky ReLU activation and batch normalization. Then, the
 816 structure is followed by a global mean pooling. After several MLP layers, the proxy outputs the cost.
 817 We use MSE loss to train this proxy. The learning rate is set as 1e-3. The batch size is set as 256.

818 **The entry-wise affine proxy:** For both f_r and g_r , we use 3 layers of GraphSAGE to encode the
 819 configuration C and the hyper-parameter α to control the DSP usage threshold with leaky Relu
 820 activation and batch normalization into the latent node representation $H \in \mathbb{R}^{|V| \times d}$ (d is the node
 821 feature dimension). The hyper-parameters of the layers are exactly the same as the proxy without
 822 constraints. Then H is separated equally into two parts $U \in \mathbb{R}^{|V| \times \frac{d}{2}}$ and $W \in \mathbb{R}^{|V| \times \frac{d}{2}}$, and we
 823 calculate $U\bar{X} + W$, after that we do the log-sum-exponential trick $\sum_{v \in V} \exp[\log(U_v \bar{X}_v + W_v) +$
 824 $\sum_{u: (u,v) \in E} \log(U_u \bar{X}_u + W_u)]$ via message passing to generate the 2-order moment entry-wise affine
 825 latent representation as introduced in Section 4. Finally, the global mean pooling and a linear layer is

826 used to obtain the output $f_r(\bar{X}; C), g_r(\bar{X}; C)$. Huber loss is used as the criterion, the learning rate is
 827 set as $5e-4$ for g_r and $1e-3$ for f_r , the batch size is 256.

828 **The entry-wise concave proxy:** The models share basically the same structure as that in the AFF
 829 proxy except for the last layers. We use linear layers with a negative Relu function to construct
 830 the CON proxy $h_r^c(\bar{X}; C) = \langle w^c, -\text{Relu}(\phi(\bar{X}; C)) \rangle + b$, as introduced in Section 4. We utilize
 831 `torch.clamp()` function to control the entries in w^c to be always greater or equal to zero in each batch
 832 of data processing during the training process.

833 **The simulated annealing baseline:** We run the simulated annealing algorithm guided by the proxy
 834 without constraints. The initial temperature is set as 1000, the cool down factor is 0.99, the ending
 835 temperature is 699. For each temperature, the number of jumps is 20. And, we set the probability for
 836 mutation is 0.1.

837 **The genetic algorithm baseline:** We run the genetic algorithm guided by the proxy without con-
 838 straints. The max generation is set as 20, the population of each generation is 40, the number of
 839 parents for mating in each generation is 20, the probability of crossover is 0.6, the probability of gene
 840 mutation is 0.01.

841 **The RL baseline:** We apply an actor critic model [57]. This model consists of 4 key components:
 842 1)States, the states are formulated as every possible partially assigned DFG; 2) Actions, given the
 843 current state and the currently considered node of the DFG, the action is whether to assign the LUT to
 844 this node; 3) State transition, given a state and an action, the probability of the next states; 4)Reward,
 845 the reward is 0 for all intermediate actions, in the last action the reward is the evaluation of the fully
 846 assigned DFG subject to the DSP usage threshold.

847 In each state at a time step, we extract the features from the last layer of the proxies without constraints
 848 f_r, g_r and concatenate them together. We utilize another GNN to encode the whole DFG into a vector
 849 encoding. The concatenated features is further combined with the vector embedding as the state
 850 encoding. Then, the state encoding is sent into the policy network that is made up of multiple MLP
 851 layers to output the critic value c and the action a which indicates whether to assign LUT for the
 852 current multiplication node. The loss for the actor is calculated by subtracting the reward by the critic
 853 value c , we use Huber loss to make c close to the reward. Note that the above scheme follows the
 854 original paper that studied the same application [30] while the status representation is based on an
 855 intermediate output given by the GNN in our proxy without constraints. The reward is defined as
 856 the negative weighted sum of LUT usage and the difference between the DSP usage and the DSP
 857 threshold:

$$r_t = \begin{cases} -\alpha f_r(X; C) - \beta \text{Relu}(t - g_r(X; C)), & s = T \\ 0, & 0 < s < T, \end{cases} \quad (22)$$

858 where T is the max step, α, β are hyper-parameters and set as 0.1, 10 respectively, t is the DSP usage
 859 threshold and s denotes the number of the step. The learning rate is set as $1e-2$, the discount factor
 860 for the reward is set as 0.95, we train the RL baseline for more than 9,000 epochs to achieve the
 861 reported performance.

862 **The mapping of $g_r(\bar{X}; C)$:** Here we introduce the mapping of the constraints in detail. The relaxed
 863 optimization goal could be written as follows:

$$\min_{\theta} f_r(\bar{X}; C), \quad \text{s.t. } g_r(\bar{X}; C) < t, \quad (23)$$

864 where $t - 1$ is the threshold for the DSP usage amount, $\bar{X} = \mathcal{A}_{\theta} \in [0, 1]^n$. As introduced in
 865 Section 2, we map the above constraints into the normalized constraint $g_r'(\bar{X}; C)$ via the following
 866 normalization.

$$g_r'(\bar{X}; C) = \frac{g_r(\bar{X}; C) - g_{\min}}{g_{\min}^+ - g_{\min}}, \quad (24)$$

867 where $g_{\min}^+ = \min_{X \in \{0,1\}^{|\mathcal{V}|} \setminus \Omega} g_r(X; C) = t$ and $g_{\min} = \min_{X \in \{0,1\}^{|\mathcal{V}|}} g_r(X; C) = 0$ in this case.
 868 Thus, the normalized constraint could be written as:

$$g_r'(\bar{X}; C) = \frac{g_r(\bar{X}; C)}{t}. \quad (25)$$

869 The constraint above could satisfy our definition of the CO problems as introduced in Section 2. The
 870 overall loss function could thus be written as follows:

$$l_r(\bar{X}; C) = f_r(\bar{X}; C) + \beta \frac{g_r(\bar{X}; C)}{t + 1}, \quad (26)$$

871 where $\beta > \max_{X \in \Omega} f(X; C)$.

872 In our implementation, we uniformly feed the network with different $\alpha = \frac{\beta}{t+1}$ for different t 's such
 873 that the model can be automatically suitable for different α 's. Simultaneously, for different t , we
 874 expect the algorithm \mathcal{A}_θ to adapt such a constraint t , so we also use t as an input, i.e., using $\mathcal{A}_\theta(\cdot; t)$.
 875 During testing, the obtained $\mathcal{A}_\theta(\cdot; t)$ outputs \bar{X} that would satisfy different DSP usage thresholds
 876 by taking different t as the input. By this, a single model could handle all ranges of DSP usage
 877 thresholds.

878 **\mathcal{A}_θ based on Gumbel-softmax trick:** We also use 3 GraphSAGE layers to encode the configuration C
 879 into the latent features. Then, we use MLP layers to reduce the dimension and map to $\mathcal{A}_\theta(C) \in [0, 1]^n$
 880 and the Gumbel-softmax trick to sample $X \sim \text{Ber}(\mathcal{A}_\theta)$. We use the soft Gumbel-softmax [48, 49]
 881 without the straight through trick. The learning rate is set as 1e-3, the batch size is 256.

882 **\mathcal{A}_θ in {relaxation (Naïve), relaxation with entry-wise concave proxy (Ours)}:** The model shares
 883 the same structure as that based on the Gumbel-softmax trick, except that the Gumbel-softmax trick
 884 is replaced by $\bar{X} \in [0, 1]^n$ directly. The learning rate is 1e-3, the batch size is 256.

885 C.5 Application III - Circuit Design for Approximate Computing

886 **Dataset details:** Each configuration C in our approximating computing (A×C) dataset is a computa-
 887 tion graph whose nodes represent either multiplication or addition calculation. For each operand, we
 888 have two different calculators to carry out the calculation: one is the precise calculator which always
 889 output the precise result but requires high computational resource workload, the other is the A×C
 890 unit which costs low computational resource but always randomly produces 10% relative error of
 891 the actual result. To balance the computation precision and the resource workload, the optimization
 892 goal is to minimize the average relative error of the computation graph given the need to use at least
 893 a certain number θ of the A×C units, where $\theta \in \{3, 5, 8\}$. For each instance in the dataset, we
 894 randomly take 1,000 different inputs to calculate the average relative error. Each input consists of 16
 895 integer numbers that are uniformly sampled from 1 to 100. To simulate the locality of some data,
 896 some of the inputs only sample 14 integers and randomly re-use two of them.

897 **The C-In, C-Out baselines:** In the C-In (C-Out) baseline, as many A×C units as the threshold
 898 requires are placed randomly near to the input (output) of the approximate computing circuit.

899 **The proxy without constraints:** We utilize 4 PNA [76] layers as the GNN backbone to show that our
 900 method is not limited with certain GNN backbones. The PNA layers take both the configuration C
 901 and the optimization variable \bar{X} as inputs with leaky Relu activation and batch normalization. Then
 902 the structure is followed by global mean pooling with MLP layers to output $f_r(\bar{X}; C)$. Huber loss is
 903 used as the criterion, the learning rate is 1e-3, and the batch size is 2048.

904 **The entry-wise affine proxy:** We also use 4 PNA layers but only take the configuration C as
 905 input to generate the latent node features $H \in \mathbb{R}^{|V| \times d}$ (d is the dimension of the node features).
 906 The hyper-parameters are exactly the same as the proxy without constraints. Then H is separated
 907 equally into two parts: $U \in \mathbb{R}^{|V| \times \frac{d}{2}}$ and $W \in \mathbb{R}^{|V| \times \frac{d}{2}}$, and we calculate $U\bar{X} + W$, after that we
 908 do the log-sum-exponential trick $\sum_{v \in V} \exp[\log(U_v \bar{X}_v + W_v) + \sum_{u: (u,v) \in E} \log(U_u \bar{X}_u + W_u)]$ via
 909 message passing to generate the 2-order moment entry-wise affine latent representation as introduced
 910 in Section 4. Huber loss is used as the criterion, the learning rate is set as 1e-3, the batch size is 2048.

911 **The entry-wise concave proxy:** The model shares basically the same structure as that in the AFF
 912 proxy except for the last layers. We utilize linear layers mixed with a $-\text{Relu}$ function to construct
 913 the CON proxy $h_r^c(\bar{X}; C) = \langle w^c, -\text{Relu}(\phi(\bar{X}; C)) \rangle + b$, as introduced in Section 4. We use
 914 torch.clamp() function to control the entries in w^c to be always greater or equal to zero in each batch
 915 of data processing during the training process.

916 **The RL baseline:** We apply an actor critic model [57]. This model consists of 4 key components:
 917 1) States, the states are formulated as every possible partially assigned $A \times C$ computation graph; 2)
 918 Actions, given the current state and the currently considered node of the $A \times C$ circuit, the action is the
 919 next node to assign with the $A \times C$ unit; 3) State transition, given a state and an action, the probability
 920 of the next states; 4) Reward, the reward is 0 for all intermediate actions, in the last action the reward
 921 is the evaluation of the fully assigned $A \times C$ computation graph.

922 In each state at a time step, we extract the features from the last layer of the proxy without constraints
 923 f_r . We utilize another GNN to encode the whole computation graph into a vector encoding. The
 924 features are further combined with the vector embedding as the state encoding. Then the state
 925 encoding is sent into the policy network that is made up of multiple MLP layers to output the critic
 926 value c and the action a which indicates the next node to assign with an $A \times C$ unit. The loss for
 927 the actor is calculated by subtracting the reward by c , and we use Huber loss to make c close to the
 928 reward. Note that the state stops if the model has already assigned with as many $A \times C$ units as the
 929 threshold requires. The reward is defined as the negative proxy prediction:

$$r_t = \begin{cases} -f_r(X; C), & s = T \\ 0, & 0 < s < T, \end{cases} \quad (27)$$

930 where T is the max step, and s denotes the number of the step. The learning rate is set as 1e-2, the
 931 discount factor for the reward is set as 0.95, we train the RL baseline for more than 9,000 epochs to
 932 achieve the reported performance.

933 **The mapping of $g_r(\bar{X}; C)$:** The relaxed optimization goal could be written as follows:

$$\min_{\theta} f_r(\bar{X}; C), \quad \text{s.t.} \quad \sum_{i=1}^n \bar{X}_i > t, \quad (28)$$

934 where $t + 1$ is the $A \times C$ unit usage threshold, $X_v = 1$ denotes the usage of an $A \times C$ unit. The
 935 relaxation of the above constraint could be written as $g'_r(\bar{X}; C) = n - \sum_{i=1}^n \bar{X}_i \in [0, n - t]$. With
 936 the method introduced in Section 2, we could normalize it as follows:

$$g_r(\bar{X}; C) = \frac{g'_r(\bar{X}; C) - g_{\min}}{g_{\min}^+ - g_{\min}}, \quad (29)$$

937 where $g_{\min}^+ = \min_{X \in \{0,1\}^n \setminus \Omega} g'_r(X; C) = n - t$ and $g_{\min} = \min_{X \in \{0,1\}^{|\mathcal{V}|}} g'_r(X; C) = 0$ in this
 938 case. Thus, the normalized constraint could be written as:

$$g_r(\bar{X}; C) = \frac{n - \sum_{i=1}^n \bar{X}_i}{n - t}. \quad (30)$$

939 The constraint above could satisfy our definition of the CO problems as introduced in Section 2, the
 940 overall loss function could thus be written as:

$$l_r(\bar{X}; C) = f_r(\bar{X}; C) + \beta \frac{n - \sum_{i=1}^n \bar{X}_i}{n - t}, \quad (31)$$

941 where $\beta > \max_{X \in \Omega} f(X; C)$.

942 In our implementation, we uniformly feed the network with different $\alpha = \frac{\beta}{n-t}$ for different t 's
 943 such that the model can be automatically suitable for different α 's. Simultaneously, for different
 944 t , we expect the algorithm \mathcal{A}_{θ} to adapt such a constraint t , so we also use t as an input, i.e., using
 945 $\mathcal{A}_{\theta}(\cdot; t)$. During testing, the obtained $\mathcal{A}_{\theta}(\cdot; t)$ outputs \bar{X} that would satisfy different $A \times C$ unit usage
 946 thresholds by taking different t as the input. By this, a single model could handle all ranges of $A \times C$
 947 unit usage thresholds.

948 **\mathcal{A}_{θ} based on Gumbel-softmax trick:** We also use 3 GraphSAGE layers with leaky Relu activation
 949 functions and batch normalization to encode the configuration C into the latent features. Then, we
 950 use MLP layers to reduce the dimension and map to $\mathcal{A}_{\theta}(C) \in [0, 1]^n$ and the Gumbel-softmax trick
 951 to sample $X \sim \text{Ber}(\mathcal{A}_{\theta})$. We use the soft Gumbel-softmax [48, 49] without the straight through trick.
 952 The learning rate is set as 1e-3, the batch size is 2048.

953 **\mathcal{A}_{θ} in {relaxation (Naïve), relaxation with entry-wise concave proxy (Ours)}:** The model shares
 954 the same structure as that based on the Gumbel-softmax trick, except that the Gumbel-softmax trick
 955 is replaced by $\bar{X} \in [0, 1]^n$ directly. The learning rate is 1e-3, the batch size is 2048.

956 **D Broader Impact**

957 In this paper, we introduce a general unsupervised framework to resolve LCO problems. The broader
958 impact of this paper is discussed from the following aspects:

959 *1) Who may benefit from this research.* The researchers, companies and organizations who utilize our
960 optimization framework to solve CO, LCO or PCO problems might benefit from this research, because
961 our framework reduces the cost of data labeling and improves the performance of the optimization.
962 In addition, more broader people might also benefit from this research, because the unsupervised
963 framework and the standardized low-cost training in comparison with the current methods mean
964 lower energy cost and less pollution, which might do good to the whole society.

965 *2) Who may be put at a risk from this research.* Although our method guarantees the quality of the
966 obtained solution when the loss is low, how much gap between the obtained solution and the optimal
967 solution is still unclear. There might be still some gaps to fill in before our method gets deployed in
968 the scenarios where rigorous approximation guarantee of the solutions is requested.

969 *3) What are the sequences of the failure of the system.* A failure of our approach will fail to give a
970 relatively good enough solution to the CO problem.

971 **E Licenses**

972 We use the following datasets in our research, their licenses are listed as follows:

- 973 • The feature based edge covering and node matching problem dataset in application I is
974 generated and proposed by us. It is inspired by [65] and utilizes the images from MNIST [66],
975 which is under the Creative Commons Attribution-Share Alike 3.0 license. The dataset is
976 publicly available.
- 977 • The resource binding problem dataset in application II is from [30] and is publicly available.
978 Please cite their paper in the new publications.
- 979 • The imprecise functional unit assignment problem dataset in application III is from [45] and
980 is publicly available. Please cite their paper in the new publications.

981 All the datasets and code bases are publicly available. They contain no human information or
982 offensive content.