

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 OPTIMIZING PSEUDO-BOOLEAN POLYNOMIALS USING HYPERGRAPH NEURAL NETWORKS

Anonymous authors

Paper under double-blind review

ABSTRACT

The challenge of solving NP-hard combinatorial optimization problems with deep learning has attracted considerable interest. Modern graph neural networks are capable of efficiently solving problem instances in an unsupervised manner, however the graph structure limits the scope of their application. We present a novel approach, hereafter referred to as PB-HGNN, which employs unsupervised Hypergraph Neural Networks (HGNNs) to solve the polynomial unconstrained binary optimization (PUBO) problem. By representing the high-order terms of the pseudo-Boolean polynomials as hyperedges in a hypergraph, the HGNN is enable to capture intricate variable interdependencies beyond pairwise interactions. As a result, our framework provides the possibility of solving a wide range of discrete problems that have not previously been addressed by neural networks. We evaluate PB-HGNN on random higher-order pseudo-Boolean polynomials, including the Sherrington-Kirkpatrick model, and max-3-SAT instances. Our results show that PB-HGNN outperforms baselines on these problems and is capable of solving large-scale PUBO instances.

1 INTRODUCTION

A pseudo-Boolean function is a mapping $f : \mathbb{B}^n \rightarrow \mathbb{R}$, where $\mathbb{B} = \{0, 1\}$ is a Boolean domain. All pseudo-Boolean functions can be uniquely represented as multilinear pseudo-Boolean polynomials Hammer & Rudeanu (1968); Boros & Hammer (2002) of the following form:

$$f(x_1, \dots, x_n) = \sum_{S \subseteq [n]} c_S \prod_{j \in S} x_j, \quad x \in \mathbb{B}^n. \quad (1)$$

The general problem to optimize (1) is known as Polynomial Unconstrained Binary Optimization (PUBO) problem. $\deg(f) = \max_{c_S \neq 0} |S|$ is the degree of pseudo-Boolean polynomial. $\deg(f) = 2$ yields the well-known Quadratic Unconstrained Binary Optimization (QUBO) problem. Recently, there has been a resurgence of interest in encoding various Combinatorial Optimization (CO) problems as QUBOs Lucas (2014); Glover et al. (2022). This is largely due to the development of specialized computational devices such as quantum annealing Boixo et al. (2013) and coherent Ising machines Wang et al. (2013), which exploit the interaction between QUBO solutions and the fundamental states of physical systems of the Ising Hamiltonian Vadlamani et al. (2020). In the last few years, new heuristic algorithms based on unsupervised graph neural networks (GNNs) were proposed for solving CO problems formulated as QUBO Schuetz et al. (2022a); Pugacheva et al. (2024); Ichikawa (2024). They encode a problem's binary variables as graph nodes and its nonzero quadratic terms as edges, and then train a GNN to minimize the continuous relaxation of the QUBO objective.

In this paper, we address the higher-order PUBO problem, when $\deg(f) \geq 3$. In combinatorial optimization, it has been widely applied for modeling a maximum satisfiability Christian et al. (2014), uncapacitated facility location Goldengorin et al. (2003); Beresnev (1973); Hammer (1968), and the p -median AlBdaiwi et al. (2009) problems. Other important applications occur in statistical mechanics Bernasconi (1987); Liers et al. (2010), computer vision Ishikawa (2011); Fix et al. (2015), protein folding Babbush et al. (2014), etc. Classical approaches for optimizing higher-order f rely on reducing the polynomial degree through quadratization methods Verma et al. (2021);

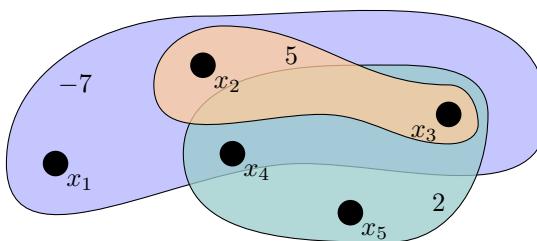


Figure 1: Hypergraph representation of $f(x_1, \dots, x_5) = 5x_2x_3 + 2x_3x_4x_5 - 7x_1x_2x_3x_4$

Chancellor et al. (2016). By introducing additional variables, the problem can be reformulated as an integer linear program or a QUBO, which can then be solved using corresponding traditional algorithms.

While in quantum and quantum-inspired optimization it is now standard to solve PUBO *directly*—without reducing to QUBO—via higher-order Ising Hamiltonians Chermoshentsev et al. (2021); Stein et al. (2023); Kanao & Goto (2022), the neural combinatorial optimization literature still predominantly relies on explicit QUBO reductions or other objectives. In this work we introduce **PB-HGNN**, a hypergraph-based unsupervised framework that operates *natively* on higher-order pseudo-Boolean polynomials: each monomial x^S becomes a hyperedge on S with weight c_S , as shown in the fig. 1, and a lightweight residual hypergraph block produces instance-specific probabilities that are refined across epochs via an epoch-level memory. By staying in the original higher-order domain, PB-HGNN preserves multi-way interactions and simplifies the modeling pipeline.

We evaluate PB-HGNN on two representative benchmarks: random higher-order PUBO instances and MAX-3-SAT (via the standard PUBO mapping) and compare against standard algorithm GNN solvers, including QUBO-based that solve problem after the quadratization procedure. Across settings, PB-HGNN attains competitive or superior solution quality while remaining simple and scalable.

Contributions.

- Direct higher-order modeling for UL. We propose PB-HGNN, the first unsupervised neural framework that solves PUBO *without* quadratization, capturing k -way structure directly.
- Memory-driven architecture. We utilize a compact epoch-level memory that carries information across iterations, stabilizing instance-specific training with minimal overhead.
- Empirical validation. On random PUBO and MAX-3-SAT, PB-HGNN matches or surpasses quadratization-based GNN baselines.

2 COMBINATORIAL OPTIMIZATION WITH UNSUPERVISED LEARNING

Unsupervised Learning (UL). Neural approaches to CO include supervised prediction Prates et al. (2019); Gasse et al. (2019) and reinforcement learning (RL) Khalil et al. (2017); Kool et al. (2019); supervised methods require labeled near-optimal solutions and often struggle to generalize, whereas RL avoids labels but is often sample-inefficient on large state spaces. A complementary route is unsupervised, instance-specific learning: optimize a differentiable model *through* which the decision variables are produced, and update model parameters to improve the instance’s objective, rather than updating the variables directly. This reparameterization often yields smoother search trajectories and fewer stalls in poor local minima Wang et al. (2022), and, by training on the instance itself, encodes useful inductive bias without requiring labeled solutions. In effect, UL turns the solver into a trainable generator of candidates that can be optimized end-to-end with standard gradient methods.

108 **GNN-based UL** Many CO problems are naturally defined on graphs, where solutions select sub-
 109 sets of nodes/edges; this makes graph neural networks appealing as the learnable maps that produce
 110 instance-specific decisions Cappart et al. (2021). Unsupervised GNN solvers include recurrent archi-
 111 tectures for constrained satisfaction problems such as RUN-CSP Tönshoff J & M (2021); instance-
 112 specific training for SAT Amizadeh et al. (2018); and Erdős Goes Neural (EGN) Karalias & Loukas
 113 (2020), which learns distributions over solutions inspired by the probabilistic method, later extended
 114 via entry-wise concave relaxations Wang et al. (2022) and meta-learning variants. Min et al. (2023)
 115 train GNN with a surrogate unsupervised loss to produce an edge ‘heat map’, and a local search then
 116 decodes a TSP tour from it. Heydaribeni et al. (2024) first proposed using hypergraph neural net-
 117 works for unsupervised combinatorial optimization. Their method, HypOp, models each constraint
 118 as a hyperedge in a constraint hypergraph and trains a HyperGNN as a learnable transformation
 119 function, with a simulated annealing step for mapping continuous relaxations back to discrete so-
 120 lutions. These methods differ in how they parametrize decisions, which unsupervised signals they
 121 optimize (constraint-violation penalties, differentiable surrogates), and how they control discretiza-
 122 tion (e.g., temperature schedules or annealing), but all share the goal of learning instance-specific
 123 heuristics directly from the objective without labeled solutions.

124 **QUBO** A prominent UL line reduces the target CO to QUBO, which is the problem to minimize
 125 a pseudo-Boolean polynomial $1 f(x)$ of degree two Hammer (1968):

$$\min_{x \in \{0,1\}^n} x^T Q x, \quad (2)$$

126 PI-GNN Schuetz et al. (2022a) replaces binary variables by probabilities $p_\theta \in [0, 1]^n$ output by
 127 a GNN and uses the relaxed QUBO as a differentiable loss, enabling general solvers whenever
 128 a QUBO formulation exists; follow-ups adapt this paradigm to specific tasks (e.g., graph color-
 129 ing Schuetz et al. (2022b); Wang et al. (2023)). To direct relaxation toward discrete solutions while
 130 keeping a smooth landscape, Ichikawa (2024) proposes a principled annealing schedule by loss
 131 regularization, that first facilitates exploration and then promotes automatic rounding. Pugacheva
 132 et al. (2024) proposed the novel GNN architecture design that feeds model outputs back into the
 133 next forward pass: Apply previous node states as dynamic node features and concatenate them with
 134 static features across epochs, improving stability and convergence of the UL procedure. Our work
 135 follows this UL paradigm but extends it to *hypergraphs* to natively capture higher-order terms in the
 136 objective function.

3 PRELIMINARIES: HYPERGRAPH NEURAL NETWORKS

140 A hypergraph $\mathcal{H} = (V, E)$ generalizes a graph by allowing each hyperedge $e \in E$ to connect
 141 an arbitrary subset of nodes $e \subseteq V$. This structure naturally captures higher-order relations that
 142 are not representable with simple pairwise edges and is widely used across machine learning and
 143 network science Kim et al. (2024); Estrada (2005). Let $|V| = n$ and $|E| = m$. We write the binary
 144 incidence matrix $A \in \{0, 1\}^{n \times m}$ with $A_{i,e} = 1$ iff $i \in e$, and define node and hyperedge degrees
 145 by $d_v(i) = \sum_e A_{i,e}$ and $d_e(e) = \sum_i A_{i,e}$, respectively. When hyperedges carry attributes (e.g.,
 146 weights, types), we denote them by a_e .

147 **Message passing on hypergraphs.** Hypergraph Neural Networks (HGNNS) extend GNN message
 148 passing by alternating *node*→*edge* and *edge*→*node* aggregations Feng et al. (2019). At layer ℓ , with
 149 node features $X^{(\ell)} \in \mathbb{R}^{n \times d_\ell}$, a general two-step update can be written as

$$(node \rightarrow edge) \quad U_e^{(\ell)} = \text{AGG}_e \{ \phi_n^{(\ell)}(X_i^{(\ell)}) : i \in e \},$$

$$(edge \rightarrow node) \quad X_i^{(\ell+1)} = \sigma \left(\psi_n^{(\ell)} \left(\bigoplus_{e \ni i} \phi_e^{(\ell)}(U_e^{(\ell)}, a_e) \right) \right),$$

150 where ϕ_n, ϕ_e, ψ_n are learnable maps, AGG_e and \oplus are permutation-invariant set aggregators (e.g.,
 151 mean/sum), and σ is an activation. This template covers a broad family of HGNNS: different normal-
 152 izations, stabilizers (e.g., BatchNorm), and choices of aggregators/MLPs lead to different concrete
 153 architectures Feng et al. (2019); Kim et al. (2024).

162 **A common instantiation (mean/sum).** A widely used choice (which we adopt further) applies
 163 linear maps before each hop and uses mean on the node→edge step and sum on the edge→node
 164 step:

$$165 \quad U_e^{(\ell)} = \frac{1}{|e|} \sum_{i \in e} W_1^{(\ell)} X_i^{(\ell)}, \quad X_i^{(\ell+1)} = \sigma \left(\text{BN} \left(\sum_{e \ni i} W_2^{(\ell)} U_e^{(\ell)} \right) \right),$$

168 with trainable $W_1^{(\ell)} \in \mathbb{R}^{d_\ell \times d_h}$, $W_2^{(\ell)} \in \mathbb{R}^{d_h \times d_{\ell+1}}$, and an optional BatchNorm (“BN”) for stability.
 169 Hyperedge attributes can modulate messages, e.g., by multiplying each $U_e^{(\ell)}$ with a scalar weight w_e
 170 or with learned embeddings Dong et al. (2020); Chien et al. (2022). Stacking two such hops with a
 171 residual (skip) connection yields a residual hypergraph block:
 172

$$173 \quad \hat{X}^{(\ell+1)} = \text{BN} \left(\sum_{e \ni i} W_2^{(\ell)} U_e^{(\ell)} \right), \quad X^{(\ell+1)} = \sigma \left(\hat{X}^{(\ell+1)} + \text{Skip}(X^{(\ell)}) \right).$$

175 Alternative normalizations based on degree matrices (e.g., variants derived from D_v and D_e) are
 176 also common and can be plugged into either hop to control scale and improve conditioning Zhou
 177 et al. (2006). The generic template above is agnostic to these choices and encompasses classical
 178 HGNNS and related approximations.

180 4 PB-HGNN: HYPERGRAPH NEURAL NETWORKS FOR PSEUDO-BOOLEAN 181 OPTIMIZATION

183 In this section, we describe a detailed explanation of the proposed framework.

185 **Hypergraph Construction.** Given a pseudo-Boolean polynomial 1 of degree $d \geq 3$, we con-
 186 struct a hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ where $\mathcal{V} = \{v_1, v_2, \dots, v_n\}$ represents the set of variables
 187 $\{x_1, x_2, \dots, x_n\}$ and $\mathcal{E} = \{e_S : S \subseteq \{1, 2, \dots, n\}, c_S \neq 0\}$ represents hyperedges, weighted
 188 by $w_S = c_S$, corresponding to polynomial terms $c_S \prod_{i \in S} x_i$. An illustrative example on a small
 189 pseudo-Boolean polynomial is shown in Figure 1. This representation preserves the exact structure
 190 of higher-order dependencies.

192 **Continuous Relaxation.** In the literature on unsupervised learning for QUBO (e.g., PI-GNN and
 193 related approaches), it is common to replace the binary decision variables $x_i \in \{0, 1\}$ by continuous
 194 probabilities $p_i \in [0, 1]$. A neural network is then trained to output these probabilities, which
 195 represent the likelihood that variable i is assigned the value 1. The resulting continuous objective is

$$197 \quad \mathcal{L}(p) = \sum_{S \subseteq [n]} c_S \prod_{i \in S} p_i, \quad p \in [0, 1]^n. \quad (3)$$

200 While this relaxation is usually motivated from a machine learning perspective, it coincides exactly
 201 with the *multilinear extension* Calinescu et al. (2007) of the underlying pseudo-Boolean function.
 202 This connection provides both an algebraic and a probabilistic interpretation of $\mathcal{L}(p)$.

203 **Lemma 1** (Multilinear extension of pseudo-Boolean polynomials). *Let $f : \{0, 1\}^n \rightarrow \mathbb{R}$ be a
 204 pseudo-Boolean polynomial 1. Then its multilinear extension coincides with $\mathcal{L}(p)$ in equation 3.*

206 The multilinear extension admits a natural probabilistic view: if $X = (X_1, \dots, X_n)$ with $X_i \sim$
 207 Bernoulli(p_i) independently, then

$$208 \quad \mathcal{L}(p) = \mathbb{E}[f(X)].$$

209 Hence the relaxed objective evaluates the expected value of the original discrete function under
 210 independent randomized rounding with marginals p . This interpretation is crucial in practice: optimiz-
 211 ing $\mathcal{L}(p)$ corresponds to searching for probability distributions over feasible solutions that yield
 212 high-quality discrete outcomes upon sampling or rounding.

214 The continuous variables are parametrized by a hypergraph neural network: $p = \sigma(\mathcal{F}_\theta(\mathcal{H}, X^{(0)}))$
 215 where σ is the sigmoid function, \mathcal{F}_θ is the HGNN with parameters θ , and $X^{(0)} \in \mathbb{R}^{n \times d_0}$ are node
 216 embeddings.

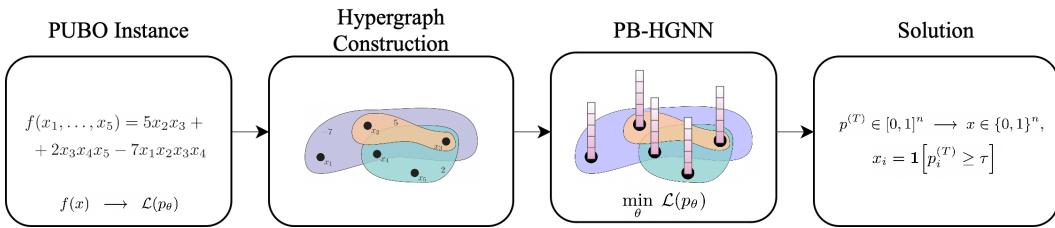


Figure 2: The PB-HGNN framework high-level design. (1) *PUBO Instance*: the input is a pseudo-Boolean polynomial and its probabilistic extension used for training. (2) *Hypergraph Construction*: each nonzero monomial becomes a hyperedge on its variables with weight equal to the coefficient. (3) *PB-HGNN*: a hypergraph neural network predicts probabilities by minimizing probabilistic extension. (4) *Solution*: After obtaining probabilities on the final iteration (epoch) T , a discrete assignment is obtained by thresholding.

Unsupervised Learning. The full PB-HGNN training framework is illustrated in Figure 3. Starting from the pseudo-Boolean polynomial, we construct the hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ as described above. Each node $v_i \in \mathcal{V}$ is initialized with an embedding $X_i^{(0)} \in \mathbb{R}^{d_0}$, drawn from a normalized random distribution or from problem-specific features if available. These embeddings are propagated through a stack of hypergraph convolutional and residual layers,

$$X^{(\ell+1)} = \phi_\ell(X^{(\ell)}, \mathcal{H}; \theta_\ell),$$

where ϕ_ℓ denotes the ℓ -th HGNN layer parameterized by θ_ℓ . After L layers, we obtain node-level hidden states $X^{(L)}$. A final projection layer followed by a sigmoid activation produces probabilities

$$p = \sigma(X^{(L)}W_{\text{out}} + \mathbf{1}b_{\text{out}}) \in [0, 1]^n, \quad X^{(L)} \in \mathbb{R}^{n \times d_L}, \quad W_{\text{out}} \in \mathbb{R}^{d_L \times 1}, \quad b_{\text{out}} \in \mathbb{R}.$$

These probabilities are interpreted as soft assignments of binary variables. The loss $\mathcal{L}(p)$ is computed directly from p and the polynomial coefficients, as detailed in the previous subsection.

Training proceeds in an unsupervised manner: given p , we evaluate the loss $\mathcal{L}(p)$, perform backpropagation through all layers of the network, and update the HGNN parameters $\theta = \{\theta_\ell, W_{\text{out}}, b_{\text{out}}\}$ using stochastic gradient descent (AdamW in our implementation). This iterative process continues until convergence or early stopping, yielding a trained PB-HGNN that maps a problem instance to a set of variable probabilities. After that, probabilities p are projected to a discrete solution by applying an indicator function,

$$x_i = \mathbf{1}\{p_i > \tau\}, \quad \tau \in (0, 1),$$

with $\tau = 0.5$ as the standard threshold. This simple rounding yields a binary assignment that can optionally be refined by local search.

4.1 PB-HGNN ARCHITECTURE

Memory-augmented hypergraph convolution. The performance of unsupervised neural solvers for combinatorial optimization depends critically on the underlying GNN architecture. In PB-HGNN we employ a memory-augmented hypergraph convolution block as the core unit. Alongside the static node embeddings $X^{(0)} \in \mathbb{R}^{n \times d_0}$, each node i carries a compact hidden memory $m_i^t \in \mathbb{R}^{d_m}$ ($d_m \ll d$), and we collect these row-wise into $M^t \in \mathbb{R}^{n \times d_m}$. At epoch t , the block takes $[X^{(0)} \parallel M^t]$ as input, performs hypergraph message passing with a residual (skip) connection, produces probabilities p^t , and updates the memory by projecting the current hidden states:

$$m_i^{t+1} = \tanh(W_m h_i^t), \quad W_m \in \mathbb{R}^{d \times d_m}.$$

where $h_i^t \in \mathbb{R}^d$ denotes the block’s output representation of node i at epoch t , and $W_m \in \mathbb{R}^{d \times d_m}$. We detach M^{t+1} from the computation graph before the next epoch, preventing backpropagation across epochs and keeping the per-epoch training cost stable. Thus the model refines its predictions over training iterations without recomputing from scratch.

270 Conceptually, our design is a lightweight variant of the iterative refinement strategy Pugacheva et al.
 271 (2024). In QI-GNN, refinement is performed across epochs by feeding the previous epoch’s pre-
 272 dicted probabilities back into the model as additional node features. PB-HGNN follows the same
 273 high-level idea to carry information across epochs, but does so via a compact hidden memory: in-
 274 stead of using raw predictions, we propagate a learned state M^t obtained from the current hidden
 275 representations and detached before the next epoch. This keeps the signal low-dimensional, re-
 276 duces noise from raw outputs, and yields a simple, stable refinement mechanism embedded in the
 277 architecture.

278

Algorithm 1 PB-HGNN: one-epoch forward/backward with epoch-level memory

280 1: **Input:** hypergraph $\mathcal{H} = (V, E)$; node features $X^{(0)} = \{x_i^{(0)}\}$; weights $\{w_e\}$; incidence
 281 $V(e), E(i)$; order emb. $\phi(\cdot)$; memory $M^t = \{m_i^t\}$
 282 2: **Output:** probabilities $P^t = \{p_i^t\}$; next memory $M^{t+1} = \{m_i^{t+1}\}$
 283 3: $h_i^{0,t} \leftarrow [x_i^{(0)} \parallel m_i^t]$; $z_i^t \leftarrow \text{ReLU}(W_{\text{in}} h_i^{0,t})$ \triangleright input fusion & projection
 284 4: $u_e^t \leftarrow \frac{1}{|e|} \sum_{i \in V(e)} W_1 z_i^t$ \triangleright node \rightarrow edge: mean aggregation
 285 5: $\tilde{a}_i^t \leftarrow \sum_{e \in E(i)} [W_2(u_e^t \odot \phi(|e|))] \cdot w_e$ \triangleright edge \rightarrow node: sum & modulation
 286 6: $a_i^t \leftarrow \text{ReLU}(\text{BN}_1(\tilde{a}_i^t))$ \triangleright norm & activation
 287 7: $u_e'^t \leftarrow \frac{1}{|e|} \sum_{i \in V(e)} W_1' a_i^t$ \triangleright second pass: node \rightarrow edge
 288 8: $\tilde{b}_i^t \leftarrow \sum_{e \in E(i)} [W_2'(u_e'^t \odot \phi(|e|))] \cdot w_e$ \triangleright second pass: edge \rightarrow node
 289 9: $b_i^t \leftarrow \text{BN}_2(\tilde{b}_i^t)$; $h_i^t \leftarrow \text{ReLU}(b_i^t + \text{Skip}(z_i^t))$ \triangleright residual (skip)
 290 10: $p_i^t \leftarrow \sigma(w_{\text{out}}^\top h_i^t + b_{\text{out}})$ \triangleright output head (sigmoid)
 291 11: $\mathcal{L} \leftarrow \sum_S c_S \prod_{j \in S} p_j^t$ \triangleright probabilistic relaxation loss
 292 12: **update** Θ with AdamW using $\nabla_{\Theta} \mathcal{L}$ \triangleright backward
 293 13: $m_i^{t+1} \leftarrow \tanh(W_m h_i^t)$; **detach** all m_i^{t+1} \triangleright memory update for next epoch
 294 14: **return** P^t, M^{t+1}

297

298

299 We use a single residual hypergraph–convolution block whose state is carried *across epochs* via
 300 a compact per-node memory. For selected epoch, we show Algorithm 1 with details of the for-
 301 ward/backword procedure. Alongside the static embeddings $X^{(0)} \in \mathbb{R}^{n \times d_0}$, each node i keeps a
 302 memory vector $m_i^t \in \mathbb{R}^{d_m}$ ($d_m \ll d$); we stack them row-wise as M^t for input fusion. In each epoch
 303 t , the block (i) concatenates features with memory and projects to the working width (line 3); (ii)
 304 performs a two-hop message pass with *mean* node \rightarrow edge aggregation (line 4) and *sum* edge \rightarrow node
 305 aggregation modulated by the hyperedge order embedding $\phi(|e|)$ and polynomial weight w_e (lines
 306 5–6); (iii) repeats the two hops with independent parameters and applies BatchNorm and a residual
 307 (skip) connection (lines 7–9). The output head produces per-node probabilities via a sigmoid (line
 308 10), and the loss is the multilinear extension of the original polynomial evaluated at these probabili-
 309 ties (line 11). We update parameters with AdamW (line 12), then update the per-node memories
 310 by a tanh projection of the current hidden states and *detach* them before the next epoch to avoid
 311 backpropagation through epochs (line 13). This realizes an iterative refinement mechanism across
 312 epochs through a lightweight, learned hidden state rather than by feeding back raw predictions.

313

314 **Training details.** By default, we use a single residual hypergraph–convolution block with Batch-
 315 Norm and ReLU, a linear output head with a sigmoid, and a compact per-node memory of size
 316 $d_m = \max\{1, \lfloor d/8 \rfloor\}$, where d is the working hidden width. Node features $X^{(0)} \in \mathbb{R}^{n \times d_0}$ are
 317 *static*: unless problem-specific features are provided, we initialize them with ℓ_2 -normalized Gaus-
 318 sian vectors. Hyperedges are stored in a sparse incidence format and carry (i) the polynomial weight
 319 $w_e = c_S$ and (ii) an order embedding $\phi(|e|)$; in our implementation we cap the order vocabulary at
 320 $|e| \leq 5$ for efficiency (higher orders are clamped). We optimize parameters θ with AdamW, apply
 321 gradient clipping (max-norm 0.5), and use a REDUCELRONPLATEAU scheduler; early stopping is
 322 triggered when the loss stabilizes. Mixed precision is enabled when available. The memory state
 323 is propagated *across epochs* and *detached* before the next epoch to avoid backpropagation through
 time. At inference, we decode by thresholding $x_i = \mathbf{1}\{p_i > \tau\}$ with default $\tau = 0.5$. For effi-
 324 ciency, we pre-compile loss terms by degree (vectorized for $d = 1$, batched for $d = 2$, tensorized

324 for $d \geq 3$). While the architecture naturally supports deeper stacks and alternative memory updates,
 325 our experiments use this minimal configuration for speed and stability.
 326

327 5 NUMERICAL EXPERIMENTS

329 5.1 RANDOM POLYNOMIALS

331 We consider synthetic pseudo-Boolean objectives of fixed degree $k \geq 3$ with n binary variables.
 332 Following prior work Chermoshentsev et al. (2021), we generate three standard classes of random
 333 instances: (i) *Class I (dense ± 1)*: each k -tuple coefficient $J_{i_1 \dots i_k}$ is sampled i.i.d. from $\{\pm 1\}$
 334 with equal probability - it corresponds to the extended version of Sherrington–Kirkpatrick model.
 335 Chermoshentsev et al. (2021). (ii) *Class II (dense uniform)*: $J_{i_1 \dots i_k} \sim \text{Unif}[-1, 1]$ i.i.d.; (iii) *Class*
 336 *III (sparse)*: start from Class II and independently drop each coefficient with probability 0.9 (set it
 337 to zero), producing a sparse higher-order interaction tensor.

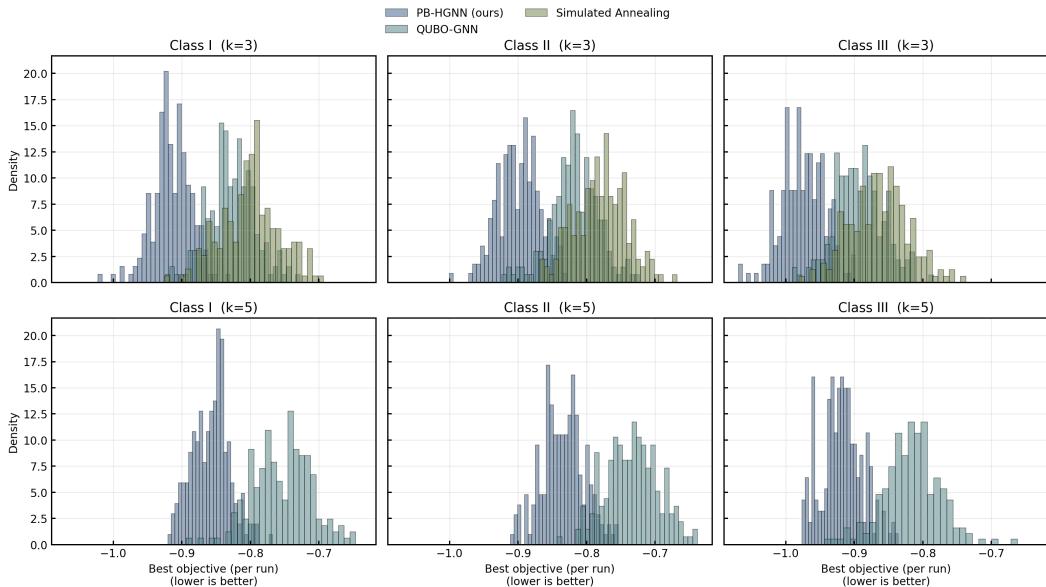
338 Let $x \in \{0, 1\}^n$. A degree- k random polynomial reads
 339

$$340 p_k(x) = \sum_{1 \leq i_1 < \dots < i_k \leq n} J_{i_1 \dots i_k} x_{i_1} \dots x_{i_k}, \quad (4)$$

342 with coefficients $J_{i_1 \dots i_k}$ drawn according to the chosen class. When needed, we also include lower-
 343 order terms to allow affine shifts:
 344

$$345 p(x) = C_\emptyset + \sum_i C_i x_i + \sum_{i < j} C_{ij} x_i x_j + \sum_{i_1 < \dots < i_k} J_{i_1 \dots i_k} x_{i_1} \dots x_{i_k}. \quad (5)$$

347 The optimization task is $\min_{x \in \{0, 1\}^n} p(x)$ (or equivalently, maximization with flipped signs). This
 348 is a native PUBO instance with hyperedges $e_{\{i_1, \dots, i_k\}}$ weighted by $J_{i_1 \dots i_k}$. ?
 349



371 Figure 3: Outcome distributions at $k = \{3, 5\}$ (PB-HGNN vs. SA). Columns correspond to the
 372 three random PUBO classes, Each panel aggregates $R=100$ independent runs (different random
 373 initializations) and shows the density of the *normalized* best objective reached within the same
 374 compute budget. We omit QUBO GNNs for $k = 5$ since it requires very large number of variables
 375 (tens thousands) in the QUBO reformulation

376 **Experimental Setup and Results.** We study random PUBO instances at fixed size $n=100$ and
 377 degrees $k \in \{3, 5\}$ for the three classes, as defined above. For each (k, class) we generate 10

378 independent instances and, unless noted, report aggregates across these instances. PB-HGNN is
 379 trained for 10,000 epochs with AdamW. We evaluate the discrete objective $p(x)$ each epoch, and
 380 keep the best-so-far value. As baselines, we use (i) a QUBO-GNN trained on a quadratized version
 381 of the same instances with the same epoch budget and optimizer, and (ii) Simulated Annealing (SA)
 382 on the PUBO form.

383 We visualize results with three figure sets. We show *Outcome distributions* 3 (2×3 grid): for
 384 each panel we aggregate the final best-of-run objectives over $R=100$ runs (multiple random ini-
 385 tializations) and display density histograms for PB-HGNN, QUBO-based GNN (that solves the
 386 QUBO problem after quadratization) : we select Pugacheva et al. (2024) as it achieves the best
 387 performance among QUBO GNNs, and Simulated Annealing, revealing both spread and tail be-
 388 havior. To make values comparable across instances with different scales and sparsities, we report
 389 $\tilde{p}(x) = p(x)/\|c\|_1$, where $\|c\|_1 = \sum_S |c_S|$ is the ℓ_1 norm of the PUBO coefficients for that in-
 390 stance; lower is better. This scale-free normalization removes instance-dependent magnitude effects
 391 (e.g., number of nonzeros, coefficient ranges), so the histograms reflect solver quality rather than
 392 raw objective scale. Across all classes, SA is competitive and lies close to PB-HGNN under this
 393 normalization, indicating that SA is not drastically worse at $k=5$. Across classes and both de-
 394 grees, PB-HGNN consistently reaches lower objectives faster than the QUBO baselines and exhibits
 395 tighter outcome distributions; the overhead plots highlight the rapidly growing variable/term counts
 396 for higher k , underscoring the advantage of operating natively on the PUBO hypergraph.

397 5.2 MAX-3-SAT

398 **Problem description.** Given Boolean variables $y_1, \dots, y_n \in \{0, 1\}$ and a 3-CNF formula with m
 399 clauses $C_r = (\ell_{r1} \vee \ell_{r2} \vee \ell_{r3})$, each literal ℓ_{rk} being either a variable y_{irk} or its negation $\neg y_{irk}$,
 400 the goal in MAX-3-SAT is to assign truth values to the variables so as to maximize the number of
 401 satisfied clauses. A clause is satisfied if at least one of its three literals is true, and violated if all
 402 three are false. The canonical objective is therefore

$$403 \max_{y \in \{0,1\}^n} \sum_{r=1}^m \mathbf{1}\{C_r(y) = \text{True}\} = m - \min_{y \in \{0,1\}^n} \sum_{r=1}^m \mathbf{1}\{C_r(y) = \text{False}\}.$$

404 We consider the equivalent minimization of the number of violated clauses (right-hand side), which
 405 is convenient for our PUBO mapping. MAX-3-SAT is NP-hard and widely used as a benchmark
 406 for discrete optimization; instances are often specified by their clause density $\alpha = m/n$ and may
 407 be *weighted* (each clause has a weight) or *unweighted* (all weights = 1). Performance is typically
 408 reported as the fraction of satisfied clauses or the number of violations at the returned assignment.

409 **PUBO formulation.** Let $x_i \in \{0, 1\}$ encode $y_i = \text{True} \iff x_i = 1$. For a literal ℓ on x_i , define
 410 its falsity factor

$$411 \bar{\ell}(x) = \begin{cases} 1 - x_i, & \ell = x_i, \\ x_i, & \ell = \neg x_i. \end{cases}$$

412 A 3-clause $C_r = (\ell_{r1} \vee \ell_{r2} \vee \ell_{r3})$ is violated iff all three literals are false, hence

$$413 v_r(x) = \prod_{k=1}^3 \bar{\ell}_{rk}(x).$$

414 Thus MAX-3-SAT becomes the unconstrained pseudo-Boolean optimization

$$415 \min_{x \in \{0,1\}^n} p(x) = \sum_{r=1}^m v_r(x) = C_\emptyset + \sum_i C_i x_i + \sum_{i < j} -C_{ij} x_i x_j + \sum_{i < j < k} C_{ijk} x_i x_j x_k,$$

416 where the coefficients $\{C_\emptyset, C_i, C_{ij}, C_{ijk}\}$ follow from expanding the products and collecting like
 417 terms. (Full coefficient formulas and derivation are provided in Appendix B.)

418 **Experimental Setup and Results.** We evaluate PB-HGNN on the MAX-3-SAT problem fol-
 419 lowing the experimental design of Yau et al. citeyau2024are. We run Two settings are consid-
 420 ered: (i) random unweighted instances with $n = 100$ variables and clause density $\alpha = m/n \in$

Table 1: Average number of unsatisfied clauses for Max-3-SAT on random instances with $N = 100$ variables and $\alpha \in \{4.00, 4.15, 4.30\}$. Standard deviation is shown after \pm . In parentheses: average runtime per instance (seconds).

Method	$\alpha = 4.00$	$\alpha = 4.15$	$\alpha = 4.30$
WalkSAT	0.94 \pm 0.92	1.46 \pm 1.11	1.97 \pm 1.28
Survey Propagation	3.32 \pm 0.81	3.87 \pm 0.79	3.94 \pm 0.93
ErdősGNN	5.46 \pm 1.91	6.14 \pm 2.01	6.79 \pm 2.03
OptGNN	4.46 \pm 1.68	5.15 \pm 1.76	5.84 \pm 2.18
PB-HGNN (ours)	3.11 \pm 0.98	4.01 \pm 1.10	4.29 \pm 1.45

4.00, 4.15, 4.30}, and (ii) benchmark instances from the UUF SATLIB dataset. In both cases, the objective is to minimize the number of violated clauses.

For the random instance setting, each formula is generated on the fly with uniformly random clauses, and results are averaged over multiple test instances. We directly replicate the baselines reported in Yau et al. (2024): *WalkSAT* (classic stochastic local search solver), *Survey Propagation* Braunstein et al. (2005), *ErdősGNN* Karalias & Loukas (2020), and unsupervised OptGNN Yau et al. (2024). We then add a row for our *PB-HGNN*. Performance is measured as the average number of unsatisfied clauses (lower is better). The results show that our method outperforms both neural solvers, while being inferior to the best classical problem-specific method.

For the second setting, we evaluate on the UUF benchmark SAT instances from SATLIB, commonly used to test incomplete Max-SAT solvers. We select 100 unsatisfiable instances from uuf125.538.100 dataset, with problems of 100 variables and 538 clauses, and report the average fraction of satisfied clauses across instances. We compare with local search sat-specific heuristics WalkSAT and Heuristic Backbone Sampling (HBS), and neural baseline *ErdősGNN* Karalias & Loukas (2020).

Table 2: Performance on UUF benchmark instances from SATLIB. We report average clause satisfaction (higher is better). PB-HGNN results will be added in the last row.

Method	average, satisfied clauses
WalkSAT	536.42
HBS	536.11
ErdősGNN	528.75
PB-HGNN (ours)	534.30

6 CONCLUSIONS

We introduced PB-HGNN, an unsupervised hypergraph–neural framework that optimizes pseudo-Boolean polynomials directly in their native higher-order form. By mapping each nonzero monomial to a hyperedge with the corresponding coefficient, and training a memory-augmented HGNN to minimize the probabilistic relaxation, our method manages to capture multi-way variable interactions. Across representative benchmark families: random higher-order PUBO, including the Sherrington–Kirkpatrick spin model, and MAX-3-SAT, PB-HGNN delivers competitive or superior solution quality to GNN baselines built on reduced quadratic formulations, heuristics, while scaling to thousands of variables. These results demonstrate that hypergraph message passing paired with instance-specific unsupervised training is a novel and effective recipe for higher-order combinatorial optimization.

486 REFERENCES
487

488 Bader F. AlBdaiwi, Diptesh Ghosh, and Boris Goldengorin. Data aggregation for p-median prob-
489 lems. *Journal of Combinatorial Optimization*, 21(3):348–363, June 2009. ISSN 1573-2886. doi:
490 10.1007/s10878-009-9251-8.

491 Saeed Amizadeh, Sergiy Matusevych, and Markus Weimer. Learning to solve circuit-sat: An un-
492 supervised differentiable approach. In *International Conference on Learning Representations*,
493 2018. URL <https://api.semanticscholar.org/CorpusID:53544639>.

494 Ryan Babbush, Alejandro Perdomo-Ortiz, Bryan O’Gorman, William Macready, and Alan Aspuru-
495 Guzik. Construction of energy functions for lattice heteropolymer models: Efficient encodings
496 for constraint satisfaction programming and quantum annealing, April 2014. ISSN 1934-4791.
497 URL <http://dx.doi.org/10.1002/9781118755815.ch05>.

498 Vladimir L. Beresnev. On a problem of mathematical standardization theory. *Upr Sist*, (5):11:43–54
499 (in Russian), 1973.

500 J. Bernasconi. Low autocorrelation binary sequences: statistical mechanics and configu-
501 ration space analysis. *Journal de Physique*, 48(4):559–567, 1987. ISSN 0302-0738.
502 doi: 10.1051/jphys:01987004804055900. URL <http://dx.doi.org/10.1051/jphys:01987004804055900>.

503 Sergio Boixo, Tameem Albash, Federico M. Spedalieri, Nicholas Chancellor, and Daniel A. Lidar.
504 Experimental signature of programmable quantum annealing. *Nature Communications*, 4(1), June
505 2013. ISSN 2041-1723. doi: 10.1038/ncomms3067. URL <http://dx.doi.org/10.1038/ncomms3067>.

506 Endre Boros and Peter L. Hammer. Pseudo-boolean optimization. *Discrete Applied Mathematics*,
507 123(1–3):155–225, November 2002. ISSN 0166-218X. doi: 10.1016/s0166-218x(01)00341-9.

508 A. Braunstein, M. Mézard, and R. Zecchina. Survey propagation: An algorithm for satisfiability.
509 *Random Structures amp; Algorithms*, 27(2):201–226, March 2005. ISSN 1098-2418. doi: 10.
510 1002/rsa.20057. URL <http://dx.doi.org/10.1002/rsa.20057>.

511 Gruiu Calinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a submodular set
512 function subject to a matroid constraint. In *IPCO*, pp. 182–196, 2007. URL https://doi.org/10.1007/978-3-540-72792-7_15.

513 Quentin Cappart, Didier Chételat, Elias B. Khalil, Andrea Lodi, Christopher Morris, and Petar
514 Veličković. Combinatorial optimization and reasoning with graph neural networks. *Proceedings
515 of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pp. 4348–
516 4355, 8 2021. doi: 10.24963/ijcai.2021/595. URL <https://doi.org/10.24963/ijcai.2021/595>.

517 N. Chancellor, S. Zohren, P. A. Warburton, S. C. Benjamin, and S. Roberts. A direct mapping of
518 max k-sat and high order parity checks to a chimera graph. *Scientific Reports*, 6(1), November
519 2016. ISSN 2045-2322. doi: 10.1038/srep37107. URL <http://dx.doi.org/10.1038/srep37107>.

520 Dmitry A. Chermoshentsev, Aleksei O. Malyshев, Mert Esencan, Egor S. Tiunov, Douglas Men-
521 doza, Alán Aspuru-Guzik, Aleksey K. Fedorov, and Alexander I. Lvovsky. Polynomial uncon-
522 strained binary optimisation inspired by optical simulation. 6 2021.

523 Eli Chien, Chao Pan, Jianhao Peng, and Olgica Milenkovic. You are allset: A multiset function
524 framework for hypergraph neural networks. In *International Conference on Learning Represen-
525 tations*, 2022. URL https://openreview.net/forum?id=hpBTIv2uy_E.

526 Kofler Christian, Greistorfer Peter, Wang Haibo, and Kochenberger Gary. A penalty function ap-
527 proach to max 3-sat problems. *Karl-Franzens-University Graz*, 2014.

528 Yihe Dong, Will Sawin, and Yoshua Bengio. Hnnh: Hypergraph networks with hyperedge
529 neurons. *ArXiv*, abs/2006.12278, 2020. URL <https://api.semanticscholar.org/CorpusID:219965680>.

540 Juan A. Estrada, Rodriguez-Velazquez. Complex networks as hypergraphs. 2005. doi: 10.48550/
 541 ARXIV.PHYSICS/0505137. URL <https://arxiv.org/abs/physics/0505137>.
 542

543 Yifan Feng, Haoxuan You, Zizhao Zhang, Rongrong Ji, and Yue Gao. Hypergraph neural net-
 544 works. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and*
 545 *Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Sym-
 546 posium on Educational Advances in Artificial Intelligence, AAAI'19/IAAI'19/EAAI'19*. AAAI
 547 Press, 2019. ISBN 978-1-57735-809-1. doi: 10.1609/aaai.v33i01.33013558. URL <https://doi.org/10.1609/aaai.v33i01.33013558>.
 548

549 Alexander Fix, Aritanan Gruber, Endre Boros, and Ramin Zabih. A hypergraph-based reduction for
 550 higher-order binary markov random fields. *IEEE Transactions on Pattern Analysis and Machine*
 551 *Intelligence*, 37(7):1387–1395, July 2015. ISSN 2160-9292. doi: 10.1109/tpami.2014.2382109.
 552 URL <http://dx.doi.org/10.1109/TPAMI.2014.2382109>.
 553

554 Maxime Gasse, Didier Chételat, Nicola Ferroni, Laurent Charlin, and Andrea Lodi. Exact combina-
 555 torial optimization with graph convolutional neural networks. In *Advances in Neural Information*
 556 *Processing Systems 32*, 2019.

557 Fred Glover, Gary Kochenberger, Rick Hennig, and Yu Du. Quantum bridge analytics i: a tutorial
 558 on formulating and using qubo models. *Annals of Operations Research*, 314(1):141–183, April
 559 2022. ISSN 1572-9338. doi: 10.1007/s10479-022-04634-2.
 560

561 Boris Goldengorin, Diptesh Ghosh, and Gerard Sierksma. Branch and peg algorithms for the simple
 562 plant location problem. *Computers & Operations Research*, 30(7):967–981, June 2003. ISSN
 563 0305-0548. doi: 10.1016/s0305-0548(02)00049-7.
 564

565 Peter L. Hammer. Plant location—a pseudo-boolean approach. *Isr J Technol*, pp. 6:330–332, 1968.
 566

567 Peter L. Hammer and Sergiu Rudeanu. *Boolean Methods in Operations Research and Related Areas*.
 568 Springer Berlin Heidelberg, 1968. ISBN 9783642858239. doi: 10.1007/978-3-642-85823-9.
 569 URL <http://dx.doi.org/10.1007/978-3-642-85823-9>.
 570

571 Nasimeh Heydaribeni, Xinrui Zhan, Ruiqi Zhang, Tina Eliassi-Rad, and Farinaz Koushanfar. Dis-
 572 tributed constrained combinatorial optimization leveraging hypergraph neural networks. *Nature*
 573 *Machine Intelligence*, 6(6):664–672, May 2024. ISSN 2522-5839. doi: 10.1038/s42256-024-00833-7. URL <http://dx.doi.org/10.1038/s42256-024-00833-7>.
 574

575 Yuma Ichikawa. Controlling continuous relaxation for combinatorial optimization, 2024.

576 H Ishikawa. Transformation of general binary mrf minimization to the first-order case. *IEEE Trans-
 577 actions on Pattern Analysis and Machine Intelligence*, 33(6):1234–1249, June 2011. ISSN 2160-
 578 9292. doi: 10.1109/tpami.2010.91. URL <http://dx.doi.org/10.1109/TPAMI.2010.91>.
 579

580 Taro Kanao and Hayato Goto. Simulated bifurcation for higher-order cost functions. *Applied Physics*
 581 *Express*, 16(1):014501, dec 2022. doi: 10.35848/1882-0786/acaba9. URL <https://dx.doi.org/10.35848/1882-0786/acaba9>.
 582

583 Nikolaos Karalias and Andreas Loukas. Erdos goes neural: an unsupervised learning framework for
 584 combinatorial optimization on graphs. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and
 585 H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 6659–6672.
 586 Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/49f85a9ed090b20c8bed85a5923c669f-Paper.pdf.
 587

588 Elias B. Khalil, Hanjun Dai, Yuyu Zhang, Bistra Dilkina, and Le Song. Learning com-
 589 binatorial optimization algorithms over graphs. In Isabelle Guyon, Ulrike von Luxburg,
 590 Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett
 591 (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neu-
 592 ral Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp.
 593 6348–6358, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/d9896106ca98d3d05b8cbdf4fd8b13a1-Abstract.html>.

594 Sunwoo Kim, Soo Yong Lee, Yue Gao, Alessia Antelmi, Mirko Polato, and Kijung Shin. A survey
 595 on hypergraph neural networks: An in-depth and step-by-step guide. In *Proceedings of the 30th*
 596 *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '24, pp. 6534–6544,
 597 New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400704901. doi:
 598 10.1145/3637528.3671457. URL <https://doi.org/10.1145/3637528.3671457>.

599 Wouter Kool, Herke van Hoof, and Max Welling. Attention, learn to solve routing problems! In
 600 *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=ByxBFsRqYm>.

601

602 Frauke Liers, Enzo Marinari, Ulrike Pagacz, Federico Ricci-Tersenghi, and Vera Schmitz. A
 603 non-disordered glassy model with a tunable interaction range. *Journal of Statistical Mechanics: Theory and Experiment*, 2010(05):L05003, May 2010. ISSN 1742-5468. doi: 10.1088/1742-5468/2010/05/L05003. URL <http://dx.doi.org/10.1088/1742-5468/2010/05/L05003>.

604

605 Andrew Lucas. Ising formulations of many np problems. *Frontiers in Physics*, 2, 2014. ISSN
 606 2296-424X. doi: 10.3389/fphy.2014.00005.

607

608 Yimeng Min, Yiwei Bai, and Carla P Gomes. Unsupervised learning for solving the travelling
 609 salesman problem. In *Thirty-seventh Conference on Neural Information Processing Systems*,
 610 2023. URL <https://openreview.net/forum?id=lAEc7aIW20>.

611

612 Marcelo Prates, Pedro H. C. Avelar, Henrique Lemos, Luis C. Lamb, and Moshe Y. Vardi. Learning
 613 to solve np-complete problems: A graph neural network for decision tsp. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):4731–4738, Jul. 2019. doi: 10.1609/aaai.v33i01.33014731. URL <https://ojs.aaai.org/index.php/AAAI/article/view/4399>.

614

615 Daria Pugacheva, Yuriy Zotov, Andrei Ermakov, and Igor Lyskov. Unsupervised graph neural
 616 networks with recurrent features for solving combinatorial optimization problems, 2024. URL
 617 <https://openreview.net/forum?id=9qtswuW5ux>.

618

619 Martin J. A. Schuetz, J. Kyle Brubaker, and Helmut G. Katzgraber. Combinatorial optimization
 620 with physics-inspired graph neural networks. *Nature Machine Intelligence*, 4(4):367–377,
 621 April 2022a. doi: 10.1038/s42256-022-00468-6. URL <https://doi.org/10.1038/s42256-022-00468-6>.

622

623 Martin J. A. Schuetz, J. Kyle Brubaker, Zhihuai Zhu, and Helmut G. Katzgraber. Graph coloring
 624 with physics-inspired graph neural networks. *Phys. Rev. Res.*, 4:043131, Nov 2022b. doi:
 625 10.1103/PhysRevResearch.4.043131. URL <https://link.aps.org/doi/10.1103/PhysRevResearch.4.043131>.

626

627 Jonas Stein, Farbod Chamanian, Maximilian Zorn, Jonas Nüßlein, Sebastian Zielinski, Michael
 628 Kölle, and Claudia Linnhoff-Popien. Evidence that pubo outperforms qubo when solving con-
 629 tinuous optimization problems with the qaoa. In *Proceedings of the Companion Conference on
 630 Genetic and Evolutionary Computation*, GECCO '23 Companion, pp. 2254–2262, New
 631 York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400701207. doi:
 632 10.1145/3583133.3596358. URL <https://doi.org/10.1145/3583133.3596358>.

633

634 Wolf H Tönshoff J, Ritzert M and Grohe M. Graph neural networks for maximum constraint sat-
 635 isfaction. *Frontiers in Artificial Intelligence*, 3, 2021. doi: 10.3389/frai.2020.580607. URL
 636 <https://doi.org/10.3389/frai.2020.580607>.

637

638 Sri Krishna Vadlamani, Tianyao Patrick Xiao, and Eli Yablonovitch. Physics successfully im-
 639 plements lagrange multiplier optimization. *Proceedings of the National Academy of Sciences*,
 640 117(43):26639–26650, October 2020. ISSN 1091-6490. doi: 10.1073/pnas.2015192117. URL
 641 <http://dx.doi.org/10.1073/pnas.2015192117>.

642

643 Amit Verma, Mark Lewis, and Gary Kochenberger. Efficient qubo transformation for higher degree
 644 pseudo boolean functions, 2021. URL <https://arxiv.org/abs/2107.11695>.

645

646 Haoyu Peter Wang, Nan Wu, Hang Yang, Cong Hao, and Pan Li. Unsupervised learning
 647 for combinatorial optimization with principled objective relaxation. 2022. URL https://openreview.net/forum?id=HjNn9oD_v47.

648 Xiangyu Wang, Xueming Yan, and Yaochu Jin. A graph neural network with negative message
 649 passing for graph coloring, 2023.
 650

651 Zhe Wang, Alireza Marandi, Kai Wen, Robert L. Byer, and Yoshihisa Yamamoto. Coherent
 652 ising machine based on degenerate optical parametric oscillators. *Phys. Rev. A*, 88:063853,
 653 Dec 2013. doi: 10.1103/PhysRevA.88.063853. URL <https://link.aps.org/doi/10.1103/PhysRevA.88.063853>.
 654

655 Morris Yau, Nikolaos Karalias, Eric Hanqing Lu, Jessica Xu, and Stefanie Jegelka. Are graph
 656 neural networks optimal approximation algorithms? In *The Thirty-eighth Annual Conference on*
 657 *Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=SxRblm9aMs>.
 658

659 Dengyong Zhou, Jiayuan Huang, and Bernhard Schölkopf. Learning with hypergraphs:
 660 Clustering, classification, and embedding. In B. Schölkopf, J. Platt, and T. Hoff-
 661 man (eds.), *Advances in Neural Information Processing Systems*, volume 19. MIT Press,
 662 2006. URL https://proceedings.neurips.cc/paper_files/paper/2006/file/dff8e9c2ac33381546d96deea9922999-Paper.pdf.
 663

665 A PROOFS

666 **Lemma 1.** (Multilinear extension of pseudo-boolean polynomial). *Let $f : \{0, 1\}^n \rightarrow \mathbb{R}$ be a
 667 pseudo-Boolean polynomial*

$$668 f(x) = \sum_{S \subseteq [n]} c_S \prod_{i \in S} x_i, \quad x \in \{0, 1\}^n. \quad (6)$$

669 Then its multilinear extension coincides with $\mathcal{L}(p)$ in equation 3, i.e.,

$$670 F(p) = \sum_{S \subseteq [n]} c_S \prod_{i \in S} p_i, \quad p \in [0, 1]^n.$$

671 *Proof.* By definition, the multilinear extension of a set function f is

$$672 F(p) = \sum_{S \subseteq [n]} f(S) \prod_{i \in S} p_i \prod_{i \notin S} (1 - p_i).$$

673 Evaluating $f(S)$ from its polynomial form gives

$$674 f(S) = f(\mathbf{1}_S) = \sum_{T \subseteq [n]} c_T \prod_{j \in T} \mathbf{1}_S(j) = \sum_{T \subseteq S} c_T.$$

675 Substituting into $F(p)$ yields

$$676 F(p) = \sum_{S \subseteq [n]} \left(\sum_{T \subseteq S} c_T \right) \prod_{i \in S} p_i \prod_{i \notin S} (1 - p_i).$$

677 Swapping the order of summation gives

$$678 F(p) = \sum_{T \subseteq [n]} c_T \sum_{S \supseteq T} \prod_{i \in S} p_i \prod_{i \notin S} (1 - p_i).$$

679 For each fixed T , factor out $\prod_{i \in T} p_i$ and write $U = S \setminus T$:

$$680 \sum_{S \supseteq T} \prod_{i \in S} p_i \prod_{i \notin S} (1 - p_i) = \left(\prod_{i \in T} p_i \right) \sum_{U \subseteq [n] \setminus T} \prod_{i \in U} p_i \prod_{i \notin U} (1 - p_i).$$

681 The inner sum is the expansion of $\prod_{i \in [n] \setminus T} (p_i + (1 - p_i)) = 1$. Hence

$$682 F(p) = \sum_{T \subseteq [n]} c_T \prod_{i \in T} p_i,$$

683 which equals $\mathcal{L}(p)$ in equation 3. □

702 A.1 HYPERPARAMETER SETTINGS
703704 Default hyperparameters optimized through extensive experiments: embedding dimension $d = 16$,
705 hidden dimension $d_h = 8$, memory dimension $d_m = 2$, number of layers $L = 2$, learning rate
706 $\alpha = 10^{-3}$ with ReduceLROnPlateau scheduling, dropout rate $\rho = 0.05$, weight decay $\lambda_2 = 10^{-6}$.
707708 A.2 EFFICIENT LOSS COMPUTATION
709710 Naively evaluating equation 3 costs $O(n^d)$. We pre-compile terms by degree:
711

712
$$\mathcal{L}(p) = \sum_{d=0}^D \mathcal{C}[d](p), \quad (7)$$

713

714 where: - $\mathcal{C}[1](p) = \sum_i c_i p_i$ (vectorized), - $\mathcal{C}[2](p) = \sum_{i,j} c_{ij} p_i p_j$ (batched matrix ops), - $\mathcal{C}[d](p) =$
715 $\sum_{S: |S|=d} c_S \prod_{i \in S} p_i$ for $d \geq 3$ (tensorized).
716717 The hypergraph is stored in sparse index format
718

719
$$\mathcal{E}_{\text{sparse}} = \{(i, e) : i \in e\},$$

720 reducing complexity from $O(nm)$ to $O(|\mathcal{E}_{\text{sparse}}|)$.
721722 B DERIVATION DETAILS FOR MAX 3 SAT PUBO
723724 For a literal on variable x_i , let the negation bit $\alpha \in \{0, 1\}$ indicate its sign ($\alpha = 0$ for x_i , $\alpha = 1$ for
725 $\neg x_i$). Define $s := 1 - 2\alpha \in \{+1, -1\}$. Then the falsity factor can be written as
726

727
$$\bar{\ell}(x) = (1 - \alpha) - s x_i.$$

728 For a clause $C_r = (\ell_{r1} \vee \ell_{r2} \vee \ell_{r3})$ on indices $I_r = \{i_{r1}, i_{r2}, i_{r3}\}$ with signs $\{\alpha_{r1}, \alpha_{r2}, \alpha_{r3}\}$ and
729 $s_{rk} = 1 - 2\alpha_{rk}$, the violation indicator expands as
730

731
$$v_r(x) = \prod_{k=1}^3 ((1 - \alpha_{rk}) - s_{rk} x_{i_{rk}}) = c_{\emptyset}^{(r)} + \sum_{p \in I_r} c_p^{(r)} x_p + \sum_{\{p, q\} \subset I_r} c_{pq}^{(r)} x_p x_q + c_{pqr}^{(r)} x_p x_q x_r,$$

732
733

734 with coefficients
735

736
$$c_{\emptyset}^{(r)} = \prod_{k=1}^3 (1 - \alpha_{rk}),$$

737
738
$$c_{i_{rk}}^{(r)} = -s_{rk} \prod_{\ell \neq k} (1 - \alpha_{r\ell}),$$

739
740
$$c_{i_{rk} i_{r\ell}}^{(r)} = s_{rk} s_{r\ell} \prod_{j \neq k, \ell} (1 - \alpha_{rj}),$$

741
742
$$c_{i_{r1} i_{r2} i_{r3}}^{(r)} = -s_{r1} s_{r2} s_{r3}.$$

743

744 Summing over clauses yields
745

746
$$p(x) = \sum_{r=1}^m v_r(x) = C_{\emptyset} + \sum_i C_i x_i + \sum_{i < j} C_{ij} x_i x_j + \sum_{i < j < k} C_{ijk} x_i x_j x_k,$$

747
748

749 where $C_{\emptyset} = \sum_r c_{\emptyset}^{(r)}$, $C_i = \sum_r c_i^{(r)}$, $C_{ij} = \sum_r c_{ij}^{(r)}$, and $C_{ijk} = \sum_r c_{ijk}^{(r)}$. Equivalently, one may
750 enumerate all 2^3 sign patterns of a clause and record the resulting monomials; both approaches are
751 algebraically identical.
752753
754
755