

# Learning to solve Minimum Cost Multicuts efficiently using Edge-Weighted Graph Convolutional Neural Networks

Steffen Jung<sup>1</sup>[0000-0001-8021-791X] (✉) and  
Margret Keuper<sup>1,2</sup>[0000-0002-8437-7993]

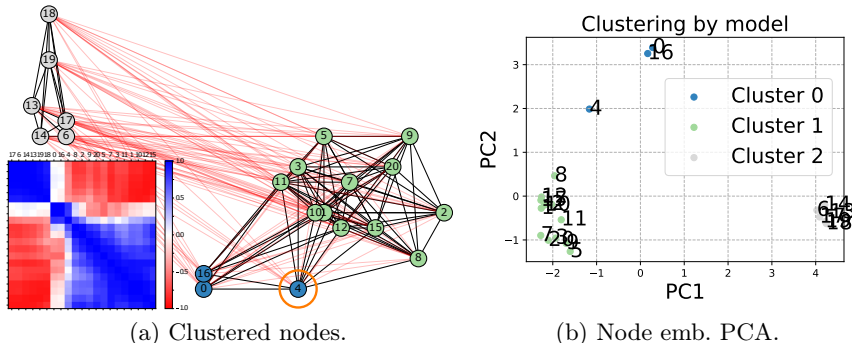
<sup>1</sup> Max Planck Institute for Informatics, Saarland Informatics Campus  
{`steffen.jung,keuper`}@mpi-inf.mpg.de  
<sup>2</sup> University of Siegen

**Abstract.** The minimum cost multicut problem is the NP-hard/APX-hard combinatorial optimization problem of partitioning a real-valued edge-weighted graph such as to minimize the total cost of the partition. While graph convolutional neural networks (GNN) have proven to be promising in the context of combinatorial optimization, most of them are only tailored to or tested on positive-valued edge weights, i.e. they do not comply with the nature of the multicut problem. We therefore adapt various GNN architectures including Graph Convolutional Networks, Signed Graph Convolutional Networks and Graph Isomorphic Networks to facilitate the efficient encoding of real-valued edge costs. Moreover, we employ a reformulation of the multicut ILP constraints to a polynomial program as loss function that allows us to learn feasible multicut solutions in a scalable way. Thus, we provide the first approach towards end-to-end trainable multicut solutions. Our findings support that GNN approaches can produce good solutions in practice while providing lower computation times and largely improved scalability compared to LP solvers and optimized heuristics, especially when considering large instances. Our code is available at <https://github.com/steffen-jung/GCN-Multicut>.

**Keywords:** Graph Neural Network · Graph Partitioning.

## 1 Introduction

Recent years have shown great advances of neural network-based approaches in various application domains from image classification [39] and natural language processing [56] up to very recent advances in decision logics [3]. While these successes indicate the importance and potential benefit of learning from data distributions, other domains such as symbolic reasoning or combinatorial optimization are still dominated by classical approaches. Recently, first attempts have been made to address specific NP-hard combinatorial problems in a learning-based setup [12, 42, 48, 50]. Specifically, such papers employ (variants of) message passing neural networks (MPNN) [20], defined on graphs [37, 45, 49] in order to model, for example, the boolean satisfiability of conjunctive normal form



**Fig. 1.** (a) Node clustering of the proposed GCN\_W\_BN on a complete graph ( $w = -220.6$ ) from IrisMP and the ordered cosine similarity between all learned node embeddings. (b) The first two principal components for each node embedding of (a). Node 4 is part of the green cluster in the optimal solution ( $w = -222.9$ ). The closeness of both solutions is reflected in the embedding.

formulas (SAT) [50] or address the travelling salesman problem [48] - both highly important NP-complete combinatorial problems. These first advances employ the ability of graph convolutional networks to efficiently learn representations of entities in graphs and prove the potential to solve hard combinatorial problems.

In this paper, we address the *minimum cost multicut* problem [5, 13], also known as the weighted correlation clustering problem (see Fig. 1). This grouping problem is substantially different from the aforementioned examples as it aims to assign binary edge labels based on a signed edge cost function. Such graph partitioning problems are ubiquitous in practical applications such as image segmentation [1, 2, 4, 31, 51], motion segmentation [30, 33], stereo matching [27], inpainting [27], object tracking [22, 32], pose tracking [24], or entity matching [47]. The minimum cost multicut problem is NP-hard, as well as APX-hard [5], which makes it a particularly challenging subject to explore. Its main difficulty lies in the exponentially growing number of constraints that define feasible solutions, especially whenever non-complete graphs are considered. Established methods solve its binary linear program formulation or linear program relaxations [27]. However, deriving optimal solutions is oftentimes intractable for large problem instances. In such cases, heuristic, iterative solvers are used as a remedy [31]. A significant disadvantage of such methods is that they can not provide gradients that would allow to train downstream tasks in an end-to-end way.

To address this issue, we propose a formulation of the minimum cost multicut problem as an MPNN. While the formulation of the multicut problem as a graph neural network seems natural, most existing GNN approaches are designed to aggregate *node* features potentially under edge constraints [53]. In contrast, instances of the multicut problem are purely defined through their *edge weights*. Graph Convolutional Networks (GCN) [37] rely on diverse node embeddings normalized by the graph Laplacian and an isotropic aggregation

function. Yet, edge weights in general and signed edge weights in particular are not modelled in standard GCNs. In this paper, we propose a simple extension of GCNs and show that the *signed* graph Laplacian can provide sufficiently strong initial node embeddings from signed edge information. This, in conjunction with an anisotropic update function which takes into account signed edge weights, facilitates GCNs to outperform more recent models such as Signed Graph Convolutional Networks (SGCN) [14], Graph Isomorphic Networks (GIN) [58] as well as models that inherently handle real-valued edge weights such as Residual Gated Graph Convolutional Networks (RGGCN) [25] and Graph Transformer Networks (GTN) [52] on the multicut problem.

To facilitate effective training, we consider a polynomial programming formulation of the minimum cost multicut problem to derive a loss function that encourages the network to issue valid solutions. Since currently available benchmarks for the minimum cost multicut problem are notoriously small, we propose two synthetic datasets with different statistics, for example w.r.t. the graph connectivity, which we use for training and analysis. We further evaluate our models on the public benchmarks BSDS300 [44], CREMI [8], and Knott3D [2].

In the following, we first briefly review the minimum cost multicut problem and commonly employed solvers. Then, we provide an overview on GNN approaches and their application in combinatorial optimization. In subsection 3.1, we present the proposed approach for solving the minimum cost multicut problem with GNNs including model adaptations and the derivation of the proposed loss function. Sec. 4 provides an empirical evaluation of the proposed approach.

## 2 The Minimum Cost Multicut Problem

The minimum cost multicut problem [11, 15] is a binary edge labelling problem defined on a graph  $G = (V, E)$ , where the connectivity is defined by edges  $e \in E \subseteq \binom{V}{2}$ , i.e.  $G$  is not necessarily complete. It allows for the definition of real-valued edge costs  $w_e \forall e \in E$ . Its solutions decompose  $G$  such as to minimize the overall cost. Specifically, the MP can be defined by the following ILP [11]:

**Definition 1.** *For a simple, connected graph  $G = (V, E)$  and an associated cost function  $w: E \rightarrow \mathbb{R}$ , written below is an instance of the multicut problem*

$$\min_{\mathbf{y} \in \{0,1\}^{|E|}} c(\mathbf{y}) = \mathbf{y}^T \mathbf{w} = \sum_{e \in E} w_e y_e \quad (1)$$

with  $y$  subject to the linear constraints

$$\forall C \in \text{cycles}(G), \forall e \in C: y_e \leq \sum_{e' \in C \setminus \{e\}} y_{e'}, \quad (2)$$

where  $\text{cycles}(\cdot)$  enumerates all cycles in graph  $G$ . The resulting  $\mathbf{y}$  is a vector of binary decision variables for each edge. Eq. (2) defines the cycle inequality constraints and ensures that if an edge is cut between two nodes, there can not be

another path in the graph connecting them. Chopra and Rao [11] further showed that the facets of the MP can be sufficiently described by cycle inequalities on all chordless cycles of  $G$ . The problem in Eq. (1)-(2) can be reformulated in a more compact way as a polynomial program (PP):

$$\min_{\mathbf{y} \in \{0,1\}^{|E|}} \sum_{e \in E} w_e y_e + K \sum_{C \in \text{cc}(G)} \sum_{e \in C} y_e \prod_{e' \in C \setminus \{e\}} (1 - y_{e'}), \quad (3)$$

for a sufficiently large penalty  $K$ . The above problem is well behaved for complete graphs where it suffices to consider all cycles of length three and Eq. (3) becomes a quadratic program. For sparse graphs, sufficient constraints may have arbitrary length  $\leq |V|$  and their enumeration might be practically infeasible. Finding an optimal solution is NP-hard and APX-hard [5]. Therefore, exact solvers are intractable for large problem instances. Linear program relaxations as well as primal feasible heuristics have been proposed to overcome this issue, which we will briefly review in the following.

*Related work on Multicut Solvers.* To solve the ILP from Definition (1), one can use general purpose LP solvers, like Gurobi [21] or CPLEX, such that optimal solutions might be in reach for small instances if the enumeration of constraints is tractable. However, no guarantees on the runtime can be provided. To mitigate the exponentially growing number of constraints, various cutting-plane [28, 34, 35] or branch-and-bound [2, 27] algorithms exist. For example, [28] employ a relaxed version of the ILP in Eq. (1) without cycle constraints. In each iteration, violated constraints are searched and added to the ILP. This approach provides optimal solutions to formerly intractable instances - yet without any runtime guarantees. Linear program relaxations [27, 34, 55] increase the tightness of the relaxation, for example using additional constraints, and provide optimality bounds. While such approaches can yield solutions within optimality bounds, their computation time can be very high and the proposed solution can be arbitrarily poor in practice. In contrast, heuristic solvers can provide runtime guarantees and have shown good results in many practical applications. The primal feasible heuristic KLj [31] iterates over pairs of partitions and computes local moves which allow to escape local optima. Competing approaches have been proposed, for example by [7, 29] or [6]. The highly efficient Greedy Additive Edge Contraction (GAEC) [31] approach aggregates nodes in a greedy procedure with an  $O(|E|\log|E|)$  worst case complexity. While such primal feasible heuristics are highly efficient in practice, they share one important draw-back with ILP solvers that becomes relevant in the learning era: they can not provide gradients that would allow for backpropagation for example to learn edge weights.

In contrast, a third order conditional random field based on the PP in Eq. (3) has been proposed in [54] and adapted in [26], which can be optimized in an end-to-end fashion using mean field iterations. This approach is strictly limited to the optimization on complete graphs. Our approach employs graph neural networks to overcome this limitation and provides a general purpose end-to-end trainable multicut approach.

### 3 Message Passing Neural Networks for Multicuts

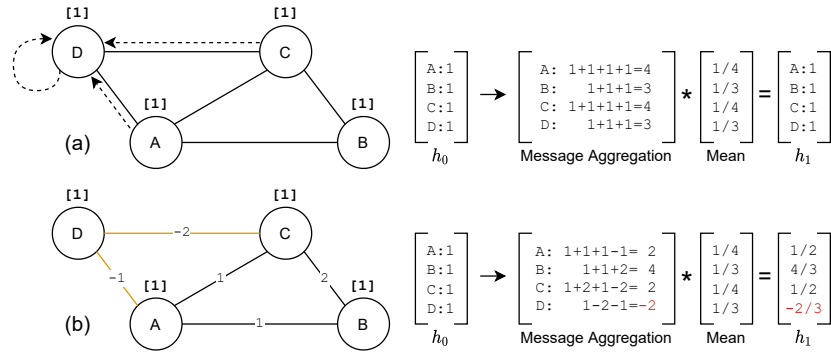
[20] provide a general framework to describe convolutions for graph data spatially as a message-passing scheme. In each convolutional layer, each node is propagating its current node features via edges to all of its neighboring nodes and updates its own features based on the messages it receives. The update is commonly described by an update function

$$\mathbf{h}_u^{(t)} = g^{(t)} \left( \mathbf{h}_u^{(t-1)}, \hat{\sum}_{v \in \mathcal{N}(u)} f^{(t)} \left( \mathbf{h}_u^{(t-1)}, \mathbf{h}_v^{(t-1)}, \mathbf{x}_{v,u} \right) \right), \quad (4)$$

where  $\mathbf{h}_u^{(t)} \in \mathbb{R}^F$  is the feature representation of node  $u$  in layer  $t$  with dimensionality  $F$ , and  $\mathbf{x}_{v,u}$  are edge features. Here,  $f$  and  $g$  are differentiable functions, and  $\hat{\Sigma}$  is a differentiable, permutation invariant aggregation function, mostly sum, max, or mean. Commonly, the message function  $f$  and the update function  $g$  are parameterized, and apply the same parameters at each location in the graph.

Various formulations have been proposed to define  $g$ . Graph Convolutional Network (GCN) [37] normalizes messages with the graph Laplacian and linearly transforms their sum to update node representations. Signed Graph Convolutional Network (SGCN) [14] aggregates messages depending on the sign of the connectivity and keeps two representations per node, one for balanced paths and one for unbalanced paths. Graph Isomorphic Network (GIN) [58] learns an injective function by defining message aggregation as a sum and learning the update function as an MLP. Residual Gated Graph Convolutional Network (RGGCN) [25] computes edge gates to aggregate messages in an anisotropic manner and learns to compute the residuals to the previous representations. Edge conditioned GCNs [53] aggregate node features using dynamic weights computed from high-dimensional edge features. Graph Transformer Network (GTN) [52] also aggregate messages anisotropically by learning a self-attention model based on the transformer models in NLP [56]. While the latter three can directly handle real-valued edge weights, all are tailored towards aggregating meaningful node features. In the following, we review recent approaches to employ such models in the context of combinatorial optimization.

*MPNNs and Combinatorial Optimization* Recently, MPNNs have been applied to several hard combinatorial optimization problems, such as the minimum vertex cover [42], maximal clique [42], maximal independent set [42], the satisfiability problem [42], and the travelling salesman problem [25]. Their objective is either to learn heuristics such as branch-and-bound variable selection policies for exact or approximate inference [16,19] or to use attention [57], reinforcement learning [9,12], or both [38,43,46] in an iterative, autoregressive procedure. [25] address the 2D Euclidean travelling salesman problem using the RGGCN model to learn edge representations. Other recent approaches address combinatorial problems by *decoding*, using supervised training such as [10]. The proposed approach is related to the work of [25], since we cast the minimum cost multicut problem as



**Fig. 2.** Message aggregation in an undirected, weighted graph where node features ( $h_0$ ) are initialized with 1. (a) Standard message aggregation in an isotropic fashion leads to no meaningful node embeddings ( $h_1 = h_0$ ). (b) Our proposed method takes edge weights into account leading to anisotropic message aggregation and meaningful node embeddings. A simple decision boundary at  $h = 0$  can now partition the graph.

a binary edge classification problem that we address using MPNN approaches, including RGGCN. We train our model in a supervised way, yet employing a dedicated loss function which encourages feasible solutions w.r.t. Eq. (2).

### 3.1 Multicut Neural Network

We cast the multicut problem into a binary edge classification task, where label  $y_{u,v} = 1$  is assigned to an edge  $(u, v)$  if it is cut, and  $y_{u,v} = 0$  otherwise. The task of the model is to learn a probability distribution  $\hat{y}_{u,v} = p(y_{u,v} = 1 | G)$  over the edges of a given graph, inferring how likely it is that an edge is cut. Based on these probabilities, we derive a configuration of edge labels,  $\mathbf{y} = \{0, 1\}^{|E|}$ . In contrast to existing autoregressive MPNN-based models in combinatorial optimization, we derive a solution after a *single forward pass* over the graph to achieve an efficient bound on the runtime of the model. In this scenario, our model can be defined by three functions, i.e.,  $\mathbf{y}_{u,v} = f_r(f_c(f_e(G, w)))$ . First,  $f_e$  is the edge representation mapping, assigning meaningful embeddings to each edge in the graph given a multicut problem instance. This function is learned by an MPNN. Second,  $f_c$  assigns to every edge its probability to be cut. This function is learned by an MLP. Last, function  $f_r$  translates the resulting configuration of edge probabilities to a feasible configuration of edge labels, hence, computes a feasible solution.

**Edge Representation Mapping** Given a multicut problem instance  $(G, w)$ , the edge representation mapping  $f_e$  learns to assign meaningful edge embeddings via MPNNs. One specific case of MPNN is GCN [37], where the node representation

update function is defined as follows:

$$\mathbf{h}_u^{(t)} = g_\theta^{(t)} \left( \mathbf{h}_u^{(t-1)} + \sum_{v \in \mathcal{N}(u)} \mathbf{L}_{[v,u]} \mathbf{h}_v^{(t-1)} \right), \quad (5)$$

where  $\mathbf{h}_u^{(t)} \in \mathbb{R}^F$  denotes the feature representation of node  $u$  in layer  $t$  with channel size  $F$ . In each layer, node representations of all neighbors of  $u$  are aggregated and normalized by  $\mathbf{L}_{[v,u]} = 1/\sqrt{\deg(u)\deg(v)}$ , where  $\mathbf{L} = \bar{\mathbf{D}}^{1/2} \tilde{\mathbf{A}} \bar{\mathbf{D}}^{1/2}$  is the normalized graph Laplacian with additional self-loops in the adjacency matrix  $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$  and degree matrix  $\bar{\mathbf{D}}$ . Conventionally,  $\mathbf{h}_u^{(0)}$  is initialized with node features  $\mathbf{x}_u$ . Intuitively, we expect normalization with the graph Laplacian to be beneficial in the MP setting, since i) its eigenvectors encode similarities of nodes within a graph [51] and ii) even sparsely connected nodes can be assigned meaningful representations [4]. However, MP instances consist of real-valued edge-weighted graphs and the normalized graph Laplacian is not defined for negative node degrees. To the best of our knowledge, there is no work enabling GCN to incorporate *real-valued* edge weights so far.

*Real-valued Edge Weights* Hence, our first task is to enable negative-valued edge weights in GCN. We can achieve this via the *signed* normalized graph Laplacian [23, 40]:

$$\bar{\mathbf{L}}_{[v,u]} = \left( \bar{\mathbf{D}}^{-1/2} \tilde{\mathbf{W}} \bar{\mathbf{D}}^{-1/2} \right)_{[v,u]} = w_{v,u} / \sqrt{\overline{\deg}(u) \overline{\deg}(v)}, \quad (6)$$

where  $\tilde{\mathbf{W}}$  is the weighted adjacency matrix and  $\bar{\mathbf{D}}$  is the signed node degree matrix with  $\overline{\deg}(u) = \sum_{v \in \mathcal{N}(u)} |w_{u,v}|$ . [18] shows that this formulation preserves the desired properties from the graph Laplacian w.r.t. encoding pairwise similarities as well as representation learning on sparsely connected nodes (see i) and ii) above).

Incorporating Eq. (6) into Eq. (5), we get

$$\mathbf{h}_u^{(t)} = g_\theta^{(t)} \left( \mathbf{h}_u^{(t-1)} + \sum_{v \in \mathcal{N}(u)} w_{v,u} \cdot (\overline{\deg}(u) \overline{\deg}(v))^{-1/2} \mathbf{h}_v^{(t-1)} \right). \quad (7)$$

Here, we can observe two new terms. First, each message is weighted by the edge weight  $w_{v,u}$  between two nodes enabling an anisotropic message-passing scheme. Fig. 2 motivates why this is necessary. While [58] show that GNNs with mean aggregation have theoretical limitations, they also note that these limitations vanish in scenarios where node features are diverse. Additionally, [58] only consider the case where neighboring nodes are aggregated in an isotropic fashion. As we show here, diverse node features are not necessary when messages are aggregated in the anisotropic fashion we propose. The resulting node representations enable distinguishing nodes in the graph despite the lack of meaningful node features. This is important in our case, since the multicut problem does not provide node features. Second, we are now able to normalize messages via the Laplacian in

real-valued graphs. The normalization acts stronger on messages that are sent to or from nodes whose adjacent edges have weights with large magnitudes. Large magnitudes on the edges usually indicate a confident decision towards joining (for positive weights) or cutting (negative weights). Thus, the normalization will allow nodes with less confident edge cues to converge to a meaningful embedding while, without such normalization, the network would notoriously focus on embedding nodes with strong edge cues, i.e. on easy decisions.

*Node Features* Conventionally, node representations at timestep 0,  $\mathbf{h}_{\mathbf{u}}^{(0)}$ , are initialized with node features  $\mathbf{x}_{\mathbf{u}}$ . However, multicut instances describe the magnitude of similarity or dissimilarity between two items via edge weights and provide no node features. Therefore, we initialize node representations with a two-dimensional vector of node degrees as:

$$\mathbf{x}_{\mathbf{u}} = \left( \sum_{v \in \mathcal{N}^+(u)} w_{u,v}, \sum_{v \in \mathcal{N}^-(u)} w_{u,v} \right), \quad (8)$$

where  $\mathcal{N}^+(u)$  is the set of neighboring nodes of  $u$  connected via positive edges, and  $\mathcal{N}^-(u)$  is the set of neighboring nodes of  $u$  connected via negative edges.

*Node-to-Edge Representation Mapping* To map two node representations to an edge representation, we use concatenation  $\mathbf{h}_{\mathbf{u},\mathbf{v}} = f_e(\mathbf{h}_{\mathbf{u}}, \mathbf{h}_{\mathbf{v}}) = \begin{pmatrix} \mathbf{h}_{\mathbf{u}} \\ \mathbf{h}_{\mathbf{v}} \end{pmatrix} \in \mathbb{R}^{2 \cdot F}$ , where  $\mathbf{h}_{\mathbf{u},\mathbf{v}}$  is the representation of edge  $(u, v)$  and  $F$  the dimension of node embedding  $\mathbf{h}_{\mathbf{u}}$ . Since we consider undirected graphs, the order of the concatenation is ambiguous. Therefore, we generate two representations for each edge, one for each direction. This doubles the number of edges to be classified in the next step. The final classification result is the average computed from both representations.

*Edge Classification* We learn edge classification function  $f_c$  via an MLP that computes likelihoods  $\hat{\mathbf{y}} \in [0, 1]^{|E|}$  for each edge in graph  $G$ , expressing the confidence whether an edge should be cut. A binary solution  $\mathbf{y} \in \{0, 1\}^{|E|}$  is retrieved by thresholding the likelihoods at 0.5. Since there is no strict guarantee that the edge label configuration  $\mathbf{y}$  is feasible w.r.t. Eq. (2), we postprocess  $\mathbf{y}$  to *round* it to feasible solutions. Therefore, we compute a connected component labelling on  $G$  after removing cut edges from  $E$  and reinstate removed edges for which both corresponding nodes remain within the same component. For efficiency, we implement the connected component labelling as a message-passing layer and can therefore assign cluster identifications to each node efficiently on the GPU.

**Training** Since we cast the multicut problem to a binary edge labelling problem, we can formulate a supervised training process that minimizes the Binary Cross Entropy (BCE) loss w.r.t. the optimal solution  $\tilde{\mathbf{y}}$ , which we denote  $\mathcal{L}_{BCE}$ .



*Cycle Consistency Loss* The BCE loss encodes feasibility only implicitly by comparison to the optimal solution. To explicitly learn feasible solutions, we take recourse to the PP formulation of the multicut problem in Eq. (3) and formulate a *feasibility loss*, that we denote *Cycle Consistency Loss* (CCL):

$$\mathcal{L}_{CCL} = \alpha \cdot \sum_{C \in cc(G,l)} \sum_{e \in C} \hat{y}_e \prod_{e' \in C \setminus \{e\}} (1 - \hat{y}_{e'}), \quad (9)$$

where  $\alpha$  is a hyperparameter, balancing BCE and CCL, and  $cc(G, l)$  is a function that returns all chordless cycles in  $G$  of length at most  $l$ . The CCL term effectively penalizes infeasible edge label configurations during training; it adds a penalty of at most  $\alpha$  for each chordless cycle that is only cut once. In practice, we only consider chordless cycles of maximum length  $l$ , and we only consider a cycle if  $e$  is cut, hence  $\hat{y}_e \geq 0.5$ . This is necessary to ensure practicable training runtimes. The total training loss is given by  $\mathcal{L} = \mathcal{L}_{BCE} + \mathcal{L}_{CCL}$ . For best results, we train all models using batch normalization.

*Training Datasets* While the multicut problem is ubiquitous in many real world applications, the amount of available annotated problem instances is scarce and domain specific. Therefore, in order to train and test a general purpose model, we generated two synthetic datasets, *IrisMP* and *RandomMP*, with complementary connectivity statistics, of 22000 multicut instances each.

The first dataset, *IrisMP*, consists of multicut instances on complete graphs based on the Iris flower dataset [17]. The generation procedure is described in the Appendix. Each problem instance consists of 120 to 276 edges. Three graphs with their respective optimal solutions are depicted in the Appendix. To complement the *IrisMP* dataset, we generated a second dataset that contains sparse but larger problem instances with 180 nodes on average, called *RandomMP*. The generation procedure is described in the Appendix. Examples are depicted in the Appendix.

## 4 Experiments

We evaluate all models trained on *IrisMP* and *RandomMP* and provide runtime as well as objective value evaluations, where we compare the proposed GCN to GIN and SGCN-based, edge-weight enabled models (see appendix for details) as well as to RGGCN [25] and GTN [52]. Then, we provide an ablation study on the proposed GCN-based edge representation mapping and the multicut loss.

### 4.1 Evaluation on Test Data

We evaluate our models on three segmentation benchmarks: a graph-based image segmentation dataset [1] based on the Berkeley Segmentation Dataset (BSDS300) [44] consisting of 100 test instances, a graph-based volume segmentation dataset [2] (Knott3D) containing 24 volumes, and 3 additional test instances based on the challenge on Circuit Reconstruction from Electron Microscopy Images

**Table 1.** Results on the test datasets. We compare different GNN variants, heuristics (GAEC) [31], LP-solver [27], and ILP-solver [27]. The performance is evaluated as optimal objective ratio  $\uparrow$  and is averaged over all datasets via harmonic mean to account for generalizability. The last column shows the total runtime  $\downarrow$  over all datasets in milliseconds. OOM indicates insufficient memory. OOT indicates no termination within 24hrs. Neither OOM nor OOT are considered in the runtime (marked with \*).

	Solver		Test Dataset						Runtime [s]		
	Variant	Depth $\alpha$ $l$	IrisMP	RandomMP	BSDS300	CREMI	Knott	h.mean	forward	total	
Proposed learned solvers	GCN_W_BN	12 0.001 3	0.9834	0.7188	<b>0.8912</b>	<b>0.7255</b>	<b>0.6902</b>	<b>0.7865</b>	0.5	4.4	
	GINO_W_BN	12 0.01 3	<b>0.9905</b>	<b>0.7387</b>	0.8474	0.5464	0.0000	0.0000	0.0	4.0	
	Signed_W_BN	12 0.01 3	0.9878	0.2526	0.6451	0.5154	0.3808	0.4510	1.3	5.3	
	RGGCN_HE	12 0.01 3	0.7976	0.1449	0.4655	0.1544	0.1735	0.2218	0.1	4.1	
	GT	12 0.001 3	0.7940	0.2964	0.6360	0.4037	0.6038	0.4836	0.1	4.0	
IrisMP	LR		0.6769	0.1118	0.6824	0.2689	0.0366	0.1164		N/A	
	MLP		0.6626	0.3127	0.7139	0.2789	0.1493	0.3051		N/A	
	RandomMP	GCN_W_BN	20 0.01 8	<b>0.9762</b>	<b>0.9041</b>	<b>0.9204</b>	0.8440	<b>0.7870</b>	<b>0.8815</b>	0.9	4.8
		GINO_W_BN	20 0.01 8	0.9528	0.8693	0.9109	0.4812	0.0000	0.0000	0.0	4.0
		Signed_W_BN	20 0.01 8	0.9709	0.8695	0.8825	0.4653	0.6408	0.7120	2.3	6.3
RGGCN_HE	20 0.01 8	0.9703	0.8787	0.8352	0.5593	OOM	-	0.1	2.7*		
Proposed learned solvers	LR		0.8035	0.3938	0.7958	<b>0.9260</b>	0.7335	0.6681		N/A	
	MLP		0.8985	0.3099	0.6804	0.4845	0.1517	0.3457		N/A	
Time-bounded solvers	GAEC		0.9836	0.9780	0.9997	0.9958	0.9968	0.9907		23.2	
	LP solver		0.3642	0.0034	0.0000	0.1516	0.0000	0.0000		6.3	
	ILP solver		0.9882	0.9525	0.9979	0.9998	OOT	-		31 918.8*	
			1.0000	1.0000	1.0000	1.0000	1.0000	1.0000		24 361.2	

(CREMI) [8] that contains volumes of electron microscopy images of fly brains. BSDS300 and Knott3D instances are available as part of a benchmark containing discrete energy minimization problems, called OpenGM [27].

**Implementation Details** We train the proposed MPNN-based solvers with (adapted) GCN, GIN, SGCN, RGGCN and GT backbones in different settings, where we uniformly set the node representation dimensionality to 128. We set the depth of the MPNN to 12 for IrisMP and 20 for RandomMP. CCL is applied with  $\alpha \in \{0, 0.01, 0.001\}$  and chordless cycle length up to 8. All of our experiments are conducted on MEGWARE Gigabyte G291-Z20 servers with NVIDIA Quadro RTX 8000 GPUs. If not stated otherwise, we consider as performance metric  $m$  the *optimal objective ratio* achieved, hence  $m = \max(0, w(\mathbf{y})/w(\tilde{\mathbf{y}})) \in [0, 1]$ .

*Results* In Tab. 1, we show the results on all test datasets of the best models based on the evaluation objective value after rounding, and thereby compare models trained on IrisMP and models trained on RandomMP. In general, sparser problems (RandomMP and established test datasets) are harder for the solvers to generalize to. This is likely due to the longer chordless cycles that the model needs to consider to ensure feasibility. Overall, our GCN-based model provides the best generalizability over all test datasets both when trained on IrisMP and RandomMP. We compare the GNN-based solvers to different baselines. First, we train logistic regression (LR) and MLPs as edge classifiers directly on the training data (concatenation of node features and edge weights). All our learned models outperform these baselines significantly. This indicates that MPNNs provide meaningful topological information to the edge classifier that facilitates solving MP instances. Second, we compare against Branch & Cut LP and ILP solvers as well as GAEC. In terms of objective value, GCN-based solvers are on par with heuristics and LP solvers on complete graphs, even when trained on sparse graphs. On general graphs, ILP solvers and GAEC issue lower energies, and, as expected, training on complete graphs does not generalize well to sparse graphs. However, the wall-clock runtime comparison shows that GCN-based solvers are faster by an order of  $10^3$  than ILP and LP solvers. They are also significantly faster than the fast and greedy GAEC heuristic. We further compared to a time-constrained version of GAEC, where we set the available time budget to the runtime of the GCN-based solver. The result shows that the trade-off between smaller energies and smaller runtime is in favor of the GCN-based solver. In the Appendix, we report additional experiments for our proposed GCN-based model on domain specific training and show that task specific priors can be learned efficiently from only a few training samples.

Next, we conduct a scalability study on random graphs with an increasing number of nodes, generated according to the RandomMP dataset. Results are shown in Tab. 2. While the GAEC is fast for small graphs, the GCN-based solver scales better and returns solutions significantly faster for larger graphs. LP and ILP solvers are not able to provide solutions within 24hrs for larger instances. It is noteworthy that GNN-based solvers spend 75-99% of their runtime rounding the solutions. Hence, GNN-based solvers are already more scalable and still have

**Table 2.** Wall-clock runtime ↓ and objective values ↓ of MPNN-based solver vs. GAEC, LP and ILP on a growing, randomly-generated graph. OOT indicates no termination within 24hrs.

Nodes	GAEC		LP		ILP		GCN_W_BN	
	[ms]	Objective	[ms]	Objective	[ms]	Objective	[ms]	Objective
$10^1$	<b>0</b>	-29	6	-24	11	-30	29	-29
$10^2$	<b>4</b>	-327	191	-246	273	-330	26	-276
$10^3$	<b>24</b>	-3051	6585	-2970	1299	-3093	29	-2643
$10^4$	228	-32 264	688 851	-31 531	18 604	-32 682	<b>78</b>	-27 552
$10^5$	2534	-323 189	OOT		2 171 134	-328 224	<b>557</b>	-269 122
$10^6$	35 181	-3 401 783	OOT		OOT		<b>8713</b>	-2 182 589

a large potential for improvement in this regard, while GAEC and LP/ILP solvers are already highly optimized for runtime.

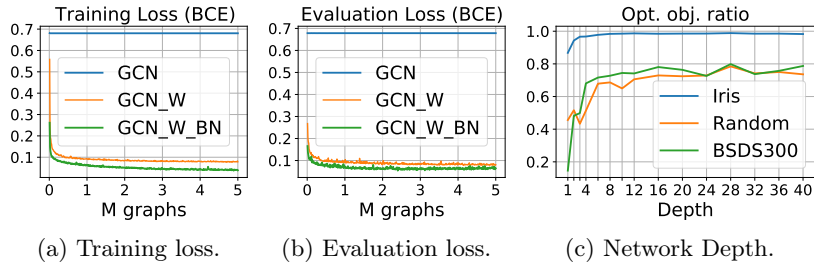
Next, we ablate on the GCN aggregation functions, loss and network depths.

*Edge-weighted GCNs* First, we determine the impact of each adjustment to the GCN update function. In Tab. 3 we show the results of this ablation study. While vanilla GCN is not applicable in the MP setting, simply removing the Laplacian from Eq. 5 provides a first baseline. We observe that adding edge weights ( $w_{u,v}$ ) to Eq. 5 improves the performance on the test split of the training data substantially. However, the model is not able to generalize to different graph statistics. By adding the signed normalization term ( $(\overline{\deg(u)\deg(v)})^{-1/2}$ ) we arrive at Eq. 7, achieving improved generalizability. Removing edge weights from Eq. 7 deteriorates performance and generalizability. Thus both changes are necessary to enable GCN in the MP setting.

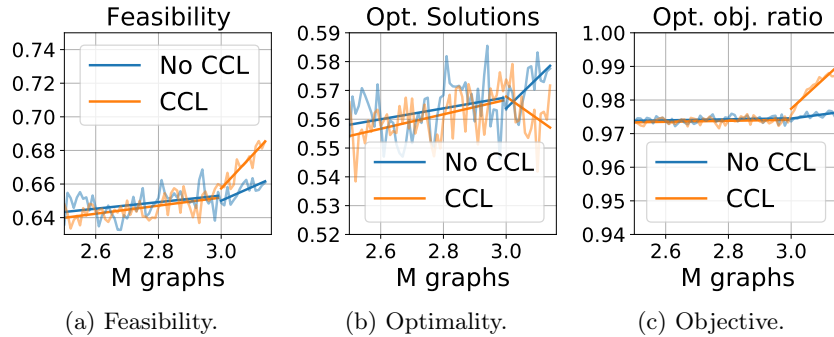
Additionally, we compare GCNs with edge weights and *signed* graph Laplacian normalization, trained with batch normalization, to the plain GCN model [37]. To this end, we train on the IrisMP dataset and set the model width to 128 and its depth to 12. Here, we set  $\alpha = 0$ , hence, we do not apply CCL. Fig. 3 shows the results of this experiment. The corresponding plots for SGCN and GIN are given in the Appendix. The variants with edge features achieve lower losses than without edge features, and batch normalization improves the loss further. In fact, the original GCN is not able to provide any meaningful features for the edge classification network. The proposed extensions enable these networks to find meaningful node representations for the multicut problem.

## 4.2 Ablation Study

*Number of Convolutional Layers* Next, we evaluate the effect of depth of the GCN model when trained on the IrisMP dataset and evaluated on IrisMP, RandomMP as well as BSDS300. Fig. 3(c) shows the results after varying the depth in increasing step sizes up to a depth of 40. The results suggest that increasing the depth improves the objective value up to a certain point. In the case of IrisMP graphs with diameter 1 and lengths of chordless cycles of at most 3, increasing



**Fig. 3.** (a) Training and (b) evaluation loss while training variants of GCN on IrisMP. Each plot compares the variants with (GCN\_W) and without (GCN) edge weights in the aggregation, and GCN\_W with batch normalization (GCN\_W\_BN). (c) Results in terms of optimal objective ratio on the evaluation data when training GCN\_W\_BN with varying depths.



**Fig. 4.** (a) Ratio of feasible solutions before repairing, (b) Ratio of optimal solutions, and (c) Optimal objective ratio, for GCN\_W\_BN on RandomMP, applying CCL after  $3M$  instances.

the depth beyond 10 has no obvious effect. This is an important observation, because [41] raise concerns that GCN models can suffer from over-smoothing such that learned representations might become indistinguishable.

*Cycle Consistency Loss* Here, we evaluate the effect of applying the cycle consistency loss from Eq. (9) by comparing models where CCL is applied after  $3M$  instances to models solely trained without CCL. Figures 4(a) and (b) show the progress of the ratio of feasible solutions and ratio of optimal solutions found during training. As soon as CCL is applied, the ratio of feasible solutions increases while the ratio of optimal solutions decreases. Hence, CCL induces a trade-off between finding feasible and optimal solutions, where the model is forced to find feasible solutions to avoid the penalty, and as a consequence, settles for suboptimal relaxed solutions. However, the objective value after rounding improves, which is most relevant because these values correspond to feasible solutions. This indicates that the model’s upper bound on the optimal energy is

**Table 3.** Ablation study with GCN [37] trained on IrisMP without CCL.. Additional comparison to vanilla versions of GIN0 [58], and MPNN [20]. We report the performance on the test data in terms of optimal objective ratio  $\uparrow$ .

Variant	IrisMP	RandMP	BSDS300	CREMI	Knott
GCN	Not applicable: Laplacian may not exist.				
- Laplacian	0.41	0.18	0.00	0.49	0.00
+ edge weights	0.95	0.18	0.40	0.57	0.19
+ signed norm.	<b>0.96</b>	<b>0.67</b>	<b>0.75</b>	<b>0.74</b>	<b>0.68</b>
= GCN_W	<b>0.96</b>	<b>0.67</b>	<b>0.75</b>	<b>0.74</b>	<b>0.68</b>
- edge weights	0.64	0.05	0.00	0.48	0.00
GIN0	0.41	0.04	0.07	0.48	0.00
MPNN	0.93	0.45	0.48	0.49	0.06

higher while the relaxation is tighter when CCL is employed. See the Appendix for an ablation on  $\alpha$  in Eq. 9.

*Meaningful Embeddings* In Fig. 1 we visualize the node embedding space given by our best performing model on an IrisMP instance. Plotting the cosine similarity between all nodes reflects the resulting clusters. This shows that the model is able to distinguish nodes based on their connectivity. We show further examples in the Appendix.

## 5 Conclusion

In this paper, we address the minimum cost multicut problem using feed forward MPNNs. To this end, we provide appropriate model and training loss modifications. Our experiments on two synthetic and two real datasets with various GCN architectures show that the proposed approach provides highly efficient solutions even to large instances and scales better than highly optimized primal feasible heuristics (GAEC), while providing competitive energies. Another significant advantage of our learning-based approach is the ability to provide gradients for downstream tasks, which we assume will inherently improve inferred solutions.

## References

1. Andres, B., Kappes, J.H., Beier, T., Köthe, U., Hamprecht, F.A.: Probabilistic Image Segmentation with Closedness Constraints. In: ICCV (2011)
2. Andres, B., Kroeger, T., Briggman, K.L., Denk, W., Korogod, N., Knott, G., Koethe, U., Hamprecht, F.A.: Globally Optimal Closed-Surface Segmentation for Connectomics. In: ECCV (2012)
3. Arakelyan, E., Daza, D., Minervini, P., Cochez, M.: Complex Query Answering with Neural Link Predictors (2021)
4. Arbeláez, P., Maire, M., Fowlkes, C., Malik, J.: Contour Detection and Hierarchical Image Segmentation. TPAMI **33**(5), 898–916 (2011)

5. Bansal, N., Blum, A., Chawla, S.: Correlation clustering. *Machine Learning* **56**(1–3), 89–113 (2004)
6. Beier, T., Andres, B., Köthe, U., Hamprecht, F.A.: An efficient fusion move algorithm for the minimum cost lifted multicut problem. In: *ECCV* (2016)
7. Beier, T., Kroeger, T., Kappes, J., Köthe, U., Hamprecht, F.: Cut, glue, & cut: A fast, approximate solver for multicut partitioning. In: *CVPR* (2014)
8. Beier, T., Pape, C., Rahaman, N., Prange, T., Berg, S., Bock, D.D., Cardona, A., Knott, G.W., Plaza, S.M., Scheffer, L.K., Koethe, U., Kreshuk, A., Hamprecht, F.A.: Multicut brings automated neurite segmentation closer to human performance. *Nature methods* **14**(2), 101–102 (2017)
9. Bello, I., Pham, H., Le, Q.V., Norouzi, M., Bengio, S.: Neural Combinatorial Optimization with Reinforcement Learning. In: *ICLR Workshop* (2017)
10. Chen, Y., Zhang, B.: Learning to solve network flow problems via neural decoding (2020)
11. Chopra, S., Rao, M.R.: The partition problem. *Mathematical Programming* **59**(1), 87–115 (1993)
12. Dai, H., Khalil, E.B., Zhang, Y., Dilkina, B., Song, L.: Learning Combinatorial Optimization Algorithms over Graphs. In: *NIPS* (2017)
13. Demaine, E.D., Emanuel, D., Fiat, A., Immorlica, N.: Correlation clustering in general weighted graphs. *Theoretical Computer Science* **361**(2–3), 172–187 (2006)
14. Derr, T., Ma, Y., Tang, J.: Signed Graph Convolutional Networks. In: *ICDM* (2018)
15. Deza, M.M., Laurent, M.: *Geometry of Cuts and Metrics*. Springer (1997)
16. Ding, J.Y., Zhang, C., Shen, L., Li, S., Wang, B., Xu, Y., Song, L.: Accelerating primal solution findings for mixed integer programs based on solution prediction. *AAAI* **34**(02), 1452–1459 (Apr 2020)
17. Fisher, R.A.: The use of multiple measurements in taxonomic problems. *Annals of Eugenics* **7**(2), 179–188 (1936)
18. Gallier, J.: Spectral theory of unsigned and signed graphs. applications to graph clustering: a survey (2016)
19. Gasse, M., Chételat, D., Ferroni, N., Charlin, L., Lodi, A.: Exact combinatorial optimization with graph convolutional neural networks. In: *NeurIPS* (2019)
20. Gilmer, J., Schoenholz, S.S., Riley, P.F., Vinyals, O., Dahl, G.E.: Neural Message Passing for Quantum Chemistry. In: *ICML* (2017)
21. Gurobi Optimization, L.: Gurobi optimizer reference manual (2020), <http://www.gurobi.com>
22. Ho, K., Kardoost, A., Pfreundt, F.J., Keuper, J., Keuper, M.: A two-stage minimum cost multicut approach to self-supervised multiple person tracking. In: *ACCV* (2020)
23. Hou, Y.P.: Bounds for the Least Laplacian Eigenvalue of a Signed Graph. *Acta Math Sinica* **21**(4), 955–960 (2005)
24. Insafutdinov, E., Pishchulin, L., Andres, B., Andriluka, M., Schiele, B.: Deepercut: A deeper, stronger, and faster multi-person pose estimation model. In: *ECCV* (2016)
25. Joshi, C.K., Laurent, T., Bresson, X.: An Efficient Graph Convolutional Network Technique for the Travelling Salesman Problem (2019)
26. Jung, S., Ziegler, S., Kardoost, A., Keuper, M.: Optimizing edge detection for image segmentation with multicut penalties. *CoRR* **abs/2112.05416** (2021)
27. Kappes, J.H., Andres, B., Hamprecht, F.A., Schnörr, C., Nowozin, S., Batra, D., Kim, S., Kausler, B.X., Kröger, T., Lellmann, J., Komodakis, N., Savchynskyy, B., Rother, C.: A Comparative Study of Modern Inference Techniques for Structured Discrete Energy Minimization Problems. *IJCV* **115**(2), 155–184 (2015)
28. Kappes, J.H., Speth, M., Andres, B., Reinelt, G., Schnörr, C.: Globally optimal image partitioning by multicuts. In: *EMMCVPR* (2011)

29. Kardoost, A., Keuper, M.: Solving minimum cost lifted multicut problems by node agglomeration. In: ACCV (2018)
30. Keuper, M., Andres, B., Brox, T.: Motion trajectory segmentation via minimum cost multicuts. In: ICCV (2015)
31. Keuper, M., Levinkov, E., Bonneel, N., Lavoué, G., Brox, T., Andres, B.: Efficient Decomposition of Image and Mesh Graphs by Lifted Multicuts. In: ICCV (2015)
32. Keuper, M., Tang, S., Andres, B., Brox, T., Schiele, B.: Motion Segmentation Multiple Object Tracking by Correlation Co-Clustering. TPAMI **42**(1), 140–153 (2020)
33. Keuper, M.: Higher-order minimum cost lifted multicuts for motion segmentation. In: ICCV (2017)
34. Kim, S., Nowozin, S., Kohli, P., Yoo, C.D.: Higher-Order Correlation Clustering for Image Segmentation. In: NIPS (2011)
35. Kim, S., Yoo, C., Nowozin, S., Kohli, P.: Image segmentation using higher-order correlation clustering. TPAMI **36**, 1761–1774 (2014)
36. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. In: ICLR (2015)
37. Kipf, T.N., Welling, M.: Semi-Supervised Classification with Graph Convolutional Networks. In: ICLR (2017)
38. Kool, W., Hoof, H.V., Welling, M.: Attention, Learn to Solve Routing Problems! In: ICLR (2019)
39. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet Classification with Deep Convolutional Neural Networks. In: NIPS (2012)
40. Kunegis, J., Schmidt, S., Lommatzsch, A., Lerner, J., Luca, E.W.D., Albayrak, S.: Spectral Analysis of Signed Graphs for Clustering, Prediction and Visualization. In: SDM (2010)
41. Li, Q., Han, Z., Wu, X.M.: Deeper Insights into Graph Convolutional Networks for Semi-Supervised Learning. In: AAAI (2018)
42. Li, Z., Chen, Q., Koltun, V.: Combinatorial Optimization with Graph Convolutional Networks and Guided Tree Search. In: NIPS (2018)
43. Ma, Q., Ge, S., He, D., Thaker, D., Drori, I.: Combinatorial Optimization by Graph Pointer Networks and Hierarchical Reinforcement Learning. In: AAAI Workshop on Deep Learning on Graphs: Methodologies and Applications (2020)
44. Martin, D., Fowlkes, C., Tal, D., Malik, J.: A Database of Human Segmented Natural Images and its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics. In: ICCV (2001)
45. Micheli, A.: Neural network for graphs: A contextual constructive approach. IEEE Transactions on Neural Networks **20**(3), 498–511 (2009)
46. Nazari, M., Oroojlooy, A., Snyder, L., Takac, M.: Reinforcement Learning for Solving the Vehicle Routing Problem. In: NIPS (2018)
47. Oulabi, Y., Bizer, C.: Extending cross-domain knowledge bases with long tail entities using web table data. In: Advances in Database Technology. pp. 385–396 (2019)
48. Prates, M.O.R., Avelar, P.H.C., Lemos, H., Lamb, L., Vardi, M.: Learning to Solve NP-Complete Problems - A Graph Neural Network for the Decision TSP. In: AAAI (2019)
49. Scarselli, F., Gori, M., Tsoi, A.C., Hagenbuchner, M., Monfardini, G.: The graph neural network model. IEEE Transactions on Neural Networks **20**(1), 61–80 (2009)
50. Selsam, D., Lamm, M., Bünz, B., Liang, P., de Moura, L., Dill, D.L.: Learning a sat solver from single-bit supervision (2019)



51. Shi, J., Malik, J.: Normalized cuts and image segmentation. *TPAMI* **22**(8), 888–905 (2000)
52. Shi, Y., Huang, Z., Wang, W., Zhong, H., Feng, S., Sun, Y.: Masked Label Prediction: Unified Message Passing Model for Semi-Supervised Classification (2020)
53. Simonovsky, M., Komodakis, N.: Dynamic edge-conditioned filters in convolutional neural networks on graphs. In: *CVPR* (2017)
54. Song, J., Andres, B., Black, M., Hilliges, O., Tang, S.: End-to-end Learning for Graph Decomposition. In: *ICCV* (2019)
55. Swoboda, P., Andres, B.: A message passing algorithm for the minimum cost multicut problem. In: *CVPR* (2017)
56. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is All You Need. In: *NIPS* (2017)
57. Vinyals, O., Fortunato, M., Jaitly, N.: Pointer Networks. In: *NIPS* (2015)
58. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How Powerful are Graph Neural Networks? In: *ICLR* (2019)