



HAL
open science

Gradient scarcity with Bilevel Optimization for Graph Learning

Hashem Ghanem, Samuel Vaiter, Nicolas Keriven

► **To cite this version:**

Hashem Ghanem, Samuel Vaiter, Nicolas Keriven. Gradient scarcity with Bilevel Optimization for Graph Learning. 2023. hal-04041721

HAL Id: hal-04041721

<https://hal.science/hal-04041721>

Preprint submitted on 22 Mar 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Gradient Scarcity with Bilevel Optimization for Graph Learning

Hashem Ghanem, Samuel Vaïter, and Nicolas Keriven

Abstract—A common issue in graph learning under the semi-supervised setting is referred to as *gradient scarcity*. That is, learning graphs by minimizing a loss on a subset of nodes causes edges between unlabelled nodes that are far from labelled ones to receive zero gradients. The phenomenon was first described when optimizing the graph and the weights of a Graph Neural Network (GNN) with a joint optimization algorithm. In this work, we give a precise mathematical characterization of this phenomenon, and prove that it also emerges in *bilevel optimization*, where additional dependency exists between the parameters of the problem. While for GNNs gradient scarcity occurs due to their finite receptive field, we show that it also occurs with the Laplacian regularization model, in the sense that gradients amplitude decreases exponentially with distance to labelled nodes. To alleviate this issue, we study several solutions: we propose to resort to latent graph learning using a Graph-to-Graph model (G2G), graph regularization to impose a prior structure on the graph, or optimizing on a larger graph than the original one with a reduced diameter. Our experiments on synthetic and real datasets validate our analysis and prove the efficiency of the proposed solutions.

Index Terms—gradient scarcity, graph learning, bilevel optimization, automatic differentiation.

I. INTRODUCTION

THE expensive cost of labelling data represents a challenge as the amount of generated data has been growing exponentially. As a result, it is common to observe both labelled and unlabelled data points, the latter being usually the vast majority. Learning tasks on datasets which comprise both labeled and unlabeled points is referred to as Semi-Supervised Learning (SSL). SSL is usually handled with extra assumptions on the data. The main one, called *homophily*, refers to the fact that “nearby” points are likely to have similar labels [1]. Moreover, points in many applications represent entities that are naturally linked to each other, *e.g.*, in biology [2] or social media [3]. There again, linked entities are likely to share the same label, which underlines the importance of exploiting the links when solving SSL inference problems. Consequently, various graph-based methods have been developed for SSL.

One issue with such methods is that their performance is highly dependent on the graph quality. This issue poses a significant challenge as real-world graphs are inherently noisy, significantly degrading the performance and leading to sub-optimal solutions. Many *graph learning* algorithms have thus been proposed in the literature to overcome this issue. Among these methods, a mainstream approach is to optimize the graph

structure by means of optimizing the performance directly in the downstream task.

This approach involves generating a graph that, when used by a graph-based model, minimizes some loss on labelled nodes. However, graph-based models themselves require an optimization process to minimize the classification loss. Therefore, both the graph learning process and the graph-based classification model need to learn by minimizing the “same” loss. There are three common gradient-based optimization routines that can be applied for this purpose. In the first routine, both the graph and the graph-based model are *jointly* optimized. In the second, *alternate minimization*, one is fixed while the other is updated in one iteration, and vice versa in the next iteration. The third routine is *bilevel optimization*, that is an (outer) graph learning optimization problem involving the optimal model obtained by an (inner) optimization.

For Graph Neural Network (GNN), the authors in [4] show that joint optimization leads to *gradient scarcity*. It refers to the fact that connections between unlabelled nodes “far” from the labelled ones receive *zero gradients*, *i.e.*, they receive no supervision during the optimization and are not learned. This is due to the finite receptive field (depth) of message-passing GNN. In this work, we focus on bilevel optimization and prove that gradient scarcity also occurs for all GNNs, despite additional dependency between the parameters. We also prove that this issue emerges with other graph-based classifiers, including Laplacian-based labels propagation, which, unlike GNNs, has infinite receptive field.

A. Semi-Supervised Learning

A graph \mathcal{G} is a pair (V, E) , where V is a set of n nodes and $E \subseteq V \times V$ is a set of edges. We represent a graph by its adjacency matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, where $A_{i,j}$ is the weight of the edge between nodes i, j . We denote by $\mathbf{X} \in \mathbb{R}^{n \times p}$ the feature matrix whose rows include the features of corresponding nodes, and by $\mathbf{Y} \in \mathbb{R}^n$ the vector of node labels.

We look at transductive SSL problems, where we have a set of points, a subset of which is labelled, and the goal is to approximate the labelling function on unlabelled points. Formally, we have $(\mathbf{X}_{obs}, \mathcal{G}_{obs}, \mathbf{Y}_{obs})$, where \mathcal{G}_{obs} is the observed graph, \mathbf{X}_{obs} are the observed node features (we will drop the subscript and write \mathbf{X} in the rest of the paper) and $\mathbf{Y}_{obs} \in \mathbb{R}^n$ contains the labels of a subset of points at coordinates $i \in V_{tr} \subset V$ and, *e.g.*, not-a-number “NaN” outside of V_{tr} . There are roughly two main strategies to solve SSL problems. The first is to *propagate* known labels using a *regularization* process. Predicted labels reads the following:

$$\mathbf{Y}_{Reg}(\mathbf{A}) \in \arg \min_{\mathbf{Y}} \frac{1}{|V_{tr}|} \sum_{i \in V_{tr}} \ell(\mathbf{Y}_i, (\mathbf{Y}_{obs})_i) + \lambda R(\mathbf{Y}, \mathbf{A}), \quad (1)$$

The authors acknowledge the support of ANR Grava ANR-18-CE40-0005 and ANR GRandMa ANR-21-CE23-0006.

HG is with CNRS, IMB, Univ. de Bourgogne; NK is with CNRS, IRISA, Univ. Rennes 1; SV is with CNRS, LJAD, Univ. Cote d’Azur.

where ℓ is a smooth loss function commonly chosen to be the Categorical Cross Entropy (CCE) loss for classification, and the Mean Square Error (MSE) for regression, R is a regularization function, and λ is a balancing parameter. A popular choice is the *Laplacian regularization*:

$$R(\mathbf{Y}, \mathbf{A}) = \frac{1}{|E|} \sum_{i,j} \mathbf{A}_{ij} (\mathbf{Y}_i - \mathbf{Y}_j)^2 = \frac{1}{|E|} \mathbf{Y}^\top \mathbf{L} \mathbf{Y}, \quad (2)$$

where $\mathbf{L} = \mathbf{L}(\mathbf{A}) = \mathbf{D} - \mathbf{A}$ is the Laplacian of the graph. Note that here the node features \mathbf{X} are not used.

The second main strategy for SSL is to train a *parametric model* $\mathbf{Y}_W(\mathbf{X}, \mathbf{A})$ parameterized by the weights W , such as GNNs. The objective reads:

$$\mathbf{Y}_{GNN}(\mathbf{A}) = \mathbf{Y}_{W^*}(\mathbf{X}, \mathbf{A}), \text{ where} \\ W^* = \arg \min_W \frac{1}{|V_{tr}|} \sum_{i \in V_{tr}} \ell \left((\mathbf{Y}_W(\mathbf{X}, \mathbf{A}))_i, (\mathbf{Y}_{obs})_i \right). \quad (3)$$

In this paper, we use message-passing GNNs with sum aggregation. The first layer is $\mathbf{X}^{[0]} = \mathbf{X}$, propagated as

$$\mathbf{X}^{[l]} = \phi \left(\mathbf{X}^{[l-1]} \mathbf{W}_1^{[l]} + \mathbf{A} \mathbf{X}^{[l-1]} \mathbf{W}_2^{[l]} + \mathbf{1}_n (\mathbf{b}^{[l]})^\top \right), \quad (4)$$

where $\mathbf{W}_1^{[l]}, \mathbf{W}_2^{[l]} \in \mathbb{R}^{d_{l-1} \times d_l}$ are learnable weights, $\mathbf{b}^{[l]} \in \mathbb{R}^{d_l}$ is a learnable bias, d_l is the output dimensionality of the l -th layer, $\mathbf{1}_n = [1, \dots, 1]^\top \in \mathbb{R}^n$, and ϕ is a non-linear function applied element-wise. The output $\mathbf{Y}_W(\mathbf{X}, \mathbf{A}) = \mathbf{X}^{[k]}$ is obtained after k rounds of message passing, and the parameters are gathered as $W = \{\mathbf{W}_1^{[l]}, \mathbf{W}_2^{[l]}, \mathbf{b}^{[l]}\}_{l=1}^k$.

B. Bilevel optimization for graph learning

We consider the case where the graph objective function is a function of the trained classifier, that is, we look at a *bilevel optimization*. Using a second set of labelled nodes $V_{out} \subset V$ distinct from V_{tr} and given a set of admissible adjacency matrices \mathcal{A} , the bilevel optimization is cast as

$$\hat{\mathbf{A}} \in \arg \min_{\mathbf{A} \in \mathcal{A}} F_{out}(\mathbf{A}) = \frac{1}{|V_{out}|} \sum_{i \in V_{out}} \ell(\mathbf{Y}(\mathbf{A})_i, (\mathbf{Y}_{obs})_i), \quad (5)$$

such that $\mathbf{Y}(\mathbf{A}) = \mathbf{Y}_{GNN}(\mathbf{A})$ or $\mathbf{Y}(\mathbf{A}) = \mathbf{Y}_{Reg}(\mathbf{A})$. That is, the minimization of the objective function F_{out} , called the *outer* optimization problem, involves $\mathbf{Y}(\mathbf{A})$, which is itself the result of an *inner* optimization problem, either (1) over \mathbf{Y} or (3) over W . Several models are possible for \mathcal{A} :

Full learning: $\mathcal{A} = [a, b]^{n \times n}$ is the set of all weighted adjacency matrices (generally with some bounds a, b on the weights). This choice necessarily leads to an impractical quadratic complexity on the minimization.

Edge refinement: the learned adjacency matrix has the same zero-pattern as the observed adjacency matrix, that is, we learn weights only on existing edges.

$$\mathcal{A} = \{\mathbf{A} \in [a, b]^{n \times n} | \mathbf{A}_{ij} = 0 \text{ when } (\mathbf{A}_{obs})_{ij} = 0\}.$$

The complexity is proportional to the number of edges, generally less than quadratic in n as graphs tend to be sparse.

Generalized edge refinement: same principle, but the zero-pattern is given by a modification of the observed adjacency matrix. For instance, taking the zero-pattern of \mathbf{A}_{obs}^r yields an

edge between neighbors that are less than r -hop from each other in \mathcal{G}_{obs} , where nodes i and j are r -hop from each other if the length of the shortest path between them in \mathcal{G}_{obs} is r .

Latent graph learning: the learned graph is the output of a parametric model, that takes as input the node features and the observed graph: $\mathcal{A} = \{\mathbf{A} = f_\theta(\mathbf{A}_{obs}, \mathbf{X})\}$. We will refer to such models as *Graph-to-Graph* (G2G).

Both the inner and the problems are treated by a gradient-based algorithms. We refer to the outer gradient ∇F_{out} , whether with respect to \mathbf{A} or θ , as *hypergradient*.

C. Contributions

Previous works observed gradient scarcity when learning the graph and a GNN classifier with joint optimization. Indeed, a k -layer GNN computes the label of a node using only information from r -hops far nodes with $r \leq k$. This label is then not a function of edges connecting nodes outside of this neighborhood, and the term in the classification loss corresponding to this label returns null gradients on those distant edges. However, it is not straightforward to extend this argument for bilevel optimization. Specifically, the previous discussion assumes that the trained weights of the GNN after gradient-based do not depend on the adjacency matrix \mathbf{A} , which is not the case in the bilevel setting. Moreover, if the problem holds in this setting, the roles of V_{tr} and V_{out} need to be clarified. Another question is if this problem is mitigated by resorting to graph-based models with infinite receptive field, e.g., the Laplacian regularization.

In this work, **we prove that hypergradient scarcity occurs under the bilevel optimization setting when adopting GNNs as a classifier**. We show that using a k -layer GNN induces null hypergradients on edges between nodes k -hop from labelled nodes in $V_{tr} \cup V_{out}$. **For the Laplacian regularization, we prove that the problem persists**, as hypergradients are exponentially damped with distance from labelled nodes. **We empirically validate our findings**. Then, we test three possible strategies to solve this issue: latent graph learning with G2G models, graph regularization and refining a power of the given adjacency matrix. Furthermore, we empirically **distinguish between hypergradient scarcity and overfitting**, in the sense that solving the former does not necessarily resolve the latter. To the best of our knowledge, this is the first work that mathematically tackles the gradient scarcity problem for bilevel optimization of graphs, and examines the phenomenon for models with infinite receptive field.

II. RELATED WORK

Bilevel optimization is used in many applications like multi-task and meta learning [5], [6], [7]. See [8] for a review of applications in different fields.

Graph learning gained in importance since real-world graphs usually have corrupted edges. The first way used to construct these graphs might be the k -nearest neighbors technique [9], [10] and its variants, but with shortcomings: we have to choose the number of nearest neighbors to consider and the associated similarity criterion. Here, we consider situations where the graph learning problem are formulated as

a *supervised* bilevel optimization problem. In [11], authors learn the parameters of Bernoulli probability distributions over independent random edges. The problem is similarly framed as a bilevel optimization, where these parameters are optimized to minimize the GNN’s validation loss. Similar to Eq. (5) with full learning of \mathbf{A} , this method includes learning n^2 parameters which limits scalability. In [12], a state-of-the-art method referred to as Graph Agreement Model (*GAM*) is proposed to learn graphs for SSL problems by penalizing the absence of an edge between nodes with the same label. Thereby the penalty is not explicitly a function of the used GNN model, and the problem isn’t bilevel. In attention mechanisms, edge weights are re-evaluated after each GNN layer based on similarity between node representations, *i.e.*, edge refinement. The similarity criterion is either user-defined like the dot product [13], [14], learned locally at each layer by a single-layer feed-forward network [15], or a combination of both schemes [16]. In contrast to bilevel optimization, these mechanisms are trained with the GNN model using joint optimization. To alleviate overfitting resulted from learning the GNN parameters and edge weights together, [17] makes use of the Label Propagation model (LPA) [18] to regularize the graph. The proposed framework produces state-of-the-art results on node classification tasks. However, authors adopt the joint optimization scenario.

Gradient scarcity was studied in [4] where the authors looked at this problem with the intuition that learning a graph in SSL problems is done to improve performance in the downstream task, thus optimizing both requires such supervision that is not available in small labelled subsets. Then, for downstream tasks adopting a k -layer GNN classifier (with $k = 2$ in their case), they identified what they refer to as the *supervision starvation problem*, which states that edges between unlabelled nodes do not receive any supervision if they are at least 2-hop from labelled nodes. They quantify the starvation for the special case of Erdős-Rényi graphs. Note that gradient scarcity and supervision starvation refer to the same phenomenon.

This issue cannot be resolved by adding more layers to the GNN as this will increase its complexity on one hand, which means more data and labels are needed, and due to the oversmoothing issue on the other hand [19]. To mitigate this issue and provide more supervision on the graph level, authors make use of the assumption that a good graph does not only perform well in labelling nodes, but also in denoising node features. Therefore, they regularize the learned graph by a contrastive loss [20], [21], [22], which evaluates its denoising performance, which overall results in a uni-level optimization.

That said, authors implicitly assumed no dependence between the GNN weights and the graph when identifying gradient scarcity, which is the case in joint/alternating optimization schemes. To the best of our knowledge, this issue has not yet been studied for the bilevel optimization setting. Moreover, it is not clear if this problem is resolved with graph-based models with infinite receptive field, *e.g.*, the Laplacian regularization. We treat both these topics in our work.

In [23], authors state that optimizing both the graph and a GNN model under the supervision of a classification task

introduces reliance on available labels, bias in the edge distribution and even reduce the span of potential application tasks. Still, this statement is not accompanied with a theoretical justification, especially regarding the first two consequences. To overcome this problem, authors suggested to avoid label-based graph optimization, and proposed an *unsupervised* graph learning framework based on contrastive learning. Although the unsupervised framework proved effective and competed state-of-the-art methods, we believe that labels contain informative knowledge that is not exploited when deploying unsupervised learners, and that better results are obtained by getting the best of both worlds.

III. HYPERGRADIENT SCARCITY WITH GNNs

In this section, we consider the bilevel optimization (5) in the *edge refinement* setting, *i.e.*, we optimize the weight of every existing edge in \mathbf{A}_{obs} , and the GNN case $\mathbf{Y}(\mathbf{A}) = \mathbf{Y}_{GNN}(\mathbf{A}) = \mathbf{Y}_{W^*}(\mathbf{A}, \mathbf{X})$.

In [4], the authors demonstrated that the predicted node label using a 2-layer GNN integrates information from nodes of distance less than two hops, *i.e.*, the label is not a function of edges connecting nodes at least 2-hop far away. Consequently, when optimizing the graph by minimizing the classification error of that label via a gradient-based algorithm, these edges receive zero-valued gradients. However, the authors used joint (or alternating) optimization of both the GNN weights and the adjacency matrix, where the dependency between W and \mathbf{A} is dropped, *i.e.*, $\mathbf{J}_W(\mathbf{A}) = \mathbf{0}$. *This is not the case for bilevel optimization.* In this section, we first examine the joint/alternating optimization schemes, and prove the existence of the problem for a generic number of layers k , similar to [4]. For the bilevel optimization setting, we then prove that the optimized weights W^* are not a function of edges connecting nodes at least k -hop from nodes in V_{tr} . After that, we conclude that hypergradient scarcity holds in the bilevel setting for edges connecting nodes at least k -hop from nodes in the *union* $V_{tr} \cup V_{out}$.

A. Scarcity for joint or alternating optimization

In this first result, we will assume that the weights W do not depend on \mathbf{A} , as is the case in joint/alternating minimization, and show gradient scarcity in $\mathbf{Y}_W(\mathbf{A}, \mathbf{X})$. This result uses the fact that W does not depend on \mathbf{A} , an hypothesis which is no longer satisfied in bilevel optimization.

Theorem III.1. *Let $\mathbf{Y}_W = \mathbf{Y}_W(\mathbf{A}, \mathbf{X})$ be the output of a k -layer GNN parameterized by W . Let i, j, u be such that nodes i, j are at least k -hop from node u . Assume that $\frac{\partial W}{\partial \mathbf{A}_{i,j}} = \mathbf{0}$. Then:*

$$\frac{\partial (\mathbf{Y}_W)_u}{\partial \mathbf{A}_{i,j}} = 0 . \quad (6)$$

Proof. The proof is done by induction on k . For $k = 1$, this is indeed the case since $\mathbf{X}^{[0]} = \mathbf{X}$ does not depend on \mathbf{A} , and that $\mathbf{A}_{i,j}$ does not belong to the row $\mathbf{A}_{u,\cdot}$, which is the only row in \mathbf{A} that contributes in the value $(\mathbf{X}^{[1]})_{u,\cdot}$.

Assume that the statement is true for some arbitrary positive integer k , we show that it is also true for a $k + 1$ -layer GNN.

If i, j are at least $(k + 1)$ -hop far from u , then clearly they are at least k -hop far from it too. Thus from the induction assumption, we have that $(\mathbf{X}^{[k]})_{u,:}$ is independent of $\mathbf{A}_{i,j}$. Also, $\mathbf{W}_1^{[k+1]}$ does not depend on $\mathbf{A}_{i,j}$ since we assume $\frac{\partial \mathbf{W}}{\partial \mathbf{A}_{i,j}} = \mathbf{0}$. Therefore, $(\mathbf{X}^{[k]} \mathbf{W}_1^{[k+1]})_{u,:}$ in (4) does not depend on $\mathbf{A}_{i,j}$ too.

In a similar way, if i, j are at least $(k + 1)$ -hop far from u , then they are at least k -hop far from any of its neighbors v where $\mathbf{A}_{u,v} \neq 0$. Therefore, if for all v , $\mathbf{A}_{u,v} \neq 0$, then $\frac{\partial (\mathbf{X}^{[k]})_{v,:}}{\partial \mathbf{A}_{i,j}} = \mathbf{0}$. Moreover, $\frac{\partial \mathbf{W}_2^{[k+1]}}{\partial \mathbf{A}_{i,j}} = \mathbf{0}$ since we assume $\frac{\partial \mathbf{W}}{\partial \mathbf{A}_{i,j}} = \mathbf{0}$. This makes $(\mathbf{A} \mathbf{X}^{[k]} \mathbf{W}_2^{[k+1]})_{u,:} = \mathbf{A}_{u,:} \mathbf{X}^{[k]} \mathbf{W}_2^{[k+1]}$ in (4) independent of $\mathbf{A}_{i,j}$. This concludes the proof, as $\frac{\partial (\mathbf{Y}_W)_u}{\partial \mathbf{A}_{i,j}} = \frac{\partial (\mathbf{X}^{[k+1]})_u}{\partial \mathbf{A}_{i,j}} = \mathbf{0}$. \square

B. Gradient of the optimized weights

Theorem III.1 assumes that W is not a function of the edge $\mathbf{A}_{i,j}$, and states, in such case, that edges between nodes at least k -hop from the training nodes used to optimize the graph (V_{out} in our case) receive no supervision. However, in the bilevel optimization scenario, after the first outer iteration W may depend on \mathbf{A} . The next theorem shows that gradient scarcity still occurs in the bilevel optimization framework, as the ‘‘optimal’’ weights used in practice are the result of a gradient-based algorithm. More precisely, we consider a sequence

$$W_{t+1} = W_t - Q_t(W_t, \nabla_{W_t} F_{in}) , \quad (7)$$

where Q_t is a smooth function. Note that W_t does not necessarily converges towards the true optimal point W^*

Theorem III.2. *Let \mathbf{A} be an input graph to a k -layer GNN with weights W , and W_t be the output obtained by optimizing (3) for W using a gradient-based iterates sequence. Let i, j be nodes that are at least k -hop from any node in V_{tr} . Then, for all $t \in \mathbb{N}$,*

$$\frac{\partial W_t(\mathbf{A})}{\partial \mathbf{A}_{i,j}} = \mathbf{0} . \quad (8)$$

Proof. The proof is carried out by induction on the iteration index t of the gradient-based optimizer. Denote by F_{in} the objective function in (3). For $t = 0$, W_0 is the initialization of W which is usually random and does not depend on \mathbf{A} . For $t \geq 0$, we assume that $\frac{\partial W_t}{\partial \mathbf{A}_{i,j}} = \mathbf{0}$ and prove this must be true for $t + 1$. By the chain rule, proving that $\frac{\partial (\nabla_{W_t} F_{in})}{\partial \mathbf{A}_{i,j}} = \mathbf{0}$ is sufficient to complete the proof. The gradient $\nabla_{W_t} F_{in}$ writes:

$$\nabla_{W_t} F_{in} = \frac{1}{|V_{tr}|} \sum_{u \in V_{tr}} \nabla_{W_t} \ell \left((\mathbf{Y}_{W_t}(\mathbf{X}, \mathbf{A}))_u, (\mathbf{Y}_{obs})_u \right) .$$

For all $u \in V_{tr}$, the term $\nabla_{W_t} \ell \left((\mathbf{Y}_{W_t}(\mathbf{X}, \mathbf{A}))_u, (\mathbf{Y}_{obs})_u \right)$ is a function of W_t and $(\mathbf{Y}_{W_t}(\mathbf{X}, \mathbf{A}))_u$. But $\frac{\partial W_t}{\partial \mathbf{A}_{i,j}} = \mathbf{0}$ from the induction assumption, and, given that, we have $\frac{\partial (\mathbf{Y}_{W_t}(\mathbf{X}, \mathbf{A}))_u}{\partial \mathbf{A}_{i,j}} = 0$ from Theorem III.1. Thus, we have for all $u \in V_{tr}$, $\frac{\partial}{\partial \mathbf{A}_{i,j}} \nabla_{W_t} \ell \left((\mathbf{Y}_{W_t}(\mathbf{X}, \mathbf{A}))_u, (\mathbf{Y}_{obs})_u \right) = \mathbf{0}$. This concludes the proof of (8) as it gives $\frac{\partial (\nabla_{W_t} F_{in})}{\partial \mathbf{A}_{i,j}} = \mathbf{0}$. \square

C. Hypergradient scarcity

Finally, we put the two previous results together. The next theorem states that within the bilevel optimization framework, edges between nodes at least k -hop from nodes in $V_{tr} \cup V_{out}$ receive no supervision.

Theorem III.3. *Let \mathbf{Y}_W be a k -layer GNN. Assume that the inner optimization problem is solved with a gradient-based algorithm (7). Then, for any pair of nodes i, j at least k -hop from nodes in $V_{out} \cup V_{tr}$, we have $\frac{\partial F_{out}}{\partial \mathbf{A}_{i,j}} = \mathbf{0}$.*

Proof. Directly from Theorem III.2 we have that $\frac{\partial W_t(\mathbf{A})}{\partial \mathbf{A}_{i,j}} = \mathbf{0}$ since i, j are at least k -hop far from nodes in V_{tr} . This makes it possible to apply Theorem III.1 to get that $\forall u \in V_{out}$; $\frac{\partial (\mathbf{Y}_{W_t})_u}{\partial \mathbf{A}_{i,j}} = \mathbf{0}$, as i, j are at least k -hop far from nodes in V_{out} and $\frac{\partial W_t(\mathbf{A})}{\partial \mathbf{A}_{i,j}} = \mathbf{0}$. This concludes the proof as F_{out} penalizes the classification error only on nodes in V_{out} . \square

Theorem III.3 shows that the hypergradient scarcity problem emerges when solving edge refinement tasks: if two nodes are at least k -hop far from nodes in $V_{out} \cup V_{tr}$ in \mathbf{A}_{obs} , the edge in between receives no hypergradients. In Section V, we will examine several strategies to mitigate this phenomenon.

IV. HYPERGRADIENT SCARCITY WITH THE LAPLACIAN REGULARIZATION

In the previous section, we have seen how the finite receptive field of GNNs directly induces the gradient scarcity problem. We now examine hypergradient scarcity when $\mathbf{Y}(\mathbf{A}) = \mathbf{Y}_{Reg}(\mathbf{A})$ with the Laplacian regularization (2). Indeed, in this case the inner problem (1) does not have a finite receptive field, in the sense that in general $\frac{\partial \mathbf{Y}(\mathbf{A})}{\partial \mathbf{A}_{i,j}} \neq 0$ for all i, j , unlike the GNN case as proven by Theorem III.1.

Surprisingly, we show that gradient scarcity still occurs in some sense. More precisely, we prove that the magnitude of hypergradients decreases exponentially with the sum of the distance to V_{tr} and the distance to V_{out} .

We consider the case where the downstream task is a regression problem, *i.e.*, ℓ in Eqs. (1) and (5) is the MSE loss function. Let $\mathbf{S}_{in} \in \mathbb{R}^{n \times n}$ be the diagonal selection matrix whose entries equal 1 if the corresponding node is in V_{tr} and 0 otherwise, the solution $\mathbf{Y}(\mathbf{A})$ enjoys a closed-form expression:

$$\mathbf{Y}(\mathbf{A}) = \left(\tilde{\mathbf{S}}_{in} + \lambda \tilde{\mathbf{L}} \right)^{-1} \tilde{\mathbf{S}}_{in} \mathbf{Y}_{obs} ,$$

where $\tilde{\mathbf{S}}_{in} = \frac{\mathbf{S}_{in}}{|V_{tr}|}$ and $\tilde{\mathbf{L}} = \frac{\mathbf{L}}{|E|}$. For simplicity from now on, we denote $\mathbf{B} = \tilde{\mathbf{S}}_{in} + \lambda \tilde{\mathbf{L}}$. Then, we write $\mathbf{Y}(\mathbf{A})$ as:

$$\mathbf{Y}(\mathbf{A}) = \mathbf{B}^{-1} \tilde{\mathbf{S}}_{in} \mathbf{Y}_{obs} . \quad (9)$$

It is well-defined thanks to the following result.

Lemma IV.1. *Assume that the graph is connected. The eigenvalues μ_i of \mathbf{B} satisfy, for all i :*

$$0 < \mu_{\min} \leq \mu_i \leq \mu_{\max} \leq \frac{1}{|V_{tr}|} + 2\lambda . \quad (10)$$

Given that, we now state the main result of this section.

Theorem IV.2. *Let nodes i, j be at least k -hop from V_{out} , and q -hop from V_{tr} . Then we have:*

$$\left| \frac{\partial F_{out}}{\partial \mathbf{A}_{ij}} \right| \lesssim \lambda \frac{\sqrt{|V_{out}|} + \mu_{\min} \sqrt{|V_{tr}|} |V_{out}|}{\mu_{\min}^3 |V_{tr}| |E|} y_{\infty}^2 (1 - \mu)^{q+k}, \quad (11)$$

where $\mu = \frac{\mu_{\min}}{\mu_{\max}}$ and $y_{\infty} = \|\mathbf{Y}_{obs}\|_{\infty}$.

Since both μ_{\min}, μ_{\max} are strictly positive, as shown in the proof in Section IV-C, then $0 < 1 - \mu < 1$. Therefore, Theorem IV.2 states that the magnitude of the hypergradient is *exponentially* damped in a speed that is at least proportional to $(1 - \mu)^{q+k}$, leading to a form of hypergradient scarcity.

The rest of this section is dedicated to proving Lemma IV.1 and Thm. IV.2. We first express $\mathbf{Y}(\mathbf{A})$ as a Neumann series, then we bound the derivative of terms in the resulted series, and by extension the gradient of F_{out} .

A. Proof of Lemma IV.1 and Neumann series expansion

In the first step, we re-write the inverse of \mathbf{B} using Neumann series. We first need to prove that $\|\mathbf{I} - \mathbf{B}\| < 1$ (see e.g., [24]), where $\mathbf{I} \in \mathbb{R}^{n \times n}$ is the identity matrix. Remark that the eigenvalues of $\mathbf{I} - \mathbf{B}$ are $1 - \mu_i$ where μ_1, \dots, μ_n are the eigenvalues of \mathbf{B} . Assuming the graph is connected, the ordered eigenvalues $\{\nu_i\}_{i=0}^n$ of $\tilde{\mathbf{L}}$ satisfy:

$$0 = \nu_1 < \nu_2 \leq \dots \leq \nu_n \leq 2. \quad (12)$$

The last inequality holds because $\|\mathbf{L}\| \leq 2d_{\max} \leq 2|E|$, where d_{\max} is the maximum degree of the graph. Let $\mathbf{u}_1, \dots, \mathbf{u}_n$ be the eigenvectors of $\tilde{\mathbf{L}}$, where $\mathbf{u}_1 \propto \mathbf{1}_n$ is associated to 0.

Proof of Lemma IV.1. We have $\|\tilde{\mathbf{S}}_{in}\| \leq 1/|V_{tr}|$ and $\|\tilde{\mathbf{L}}\| \leq 2$ so by a triangular inequality the upper bound is proved.

Using the eigendecomposition of $\tilde{\mathbf{L}}$ and recalling that $\nu_1 = 0$, for any $\mathbf{x} \in \mathbb{R}^n$:

$$\begin{aligned} \mathbf{x}^{\top} \mathbf{B} \mathbf{x} &= \lambda \mathbf{x}^{\top} \tilde{\mathbf{L}} \mathbf{x} + \mathbf{x}^{\top} \tilde{\mathbf{S}}_{in} \mathbf{x} \\ &= \lambda \sum_{i=2}^n (\mathbf{x}^{\top} \mathbf{u}_i)^2 \nu_i + \frac{\sum_{i \in V_{tr}} \mathbf{x}_i^2}{|V_{tr}|} \end{aligned}$$

which, minimized over the unit sphere, gives the expression of μ_{\min} . It is immediate that $\mu_{\min} \geq 0$. We prove that this value is strictly positive. Indeed, $\mathbf{x}^{\top} \mathbf{B} \mathbf{x} = 0$ implies that $\mathbf{x}^{\top} \tilde{\mathbf{S}}_{in} \mathbf{x} = 0$ and therefore $\mathbf{x}_i = 0$ for $i \in V_{tr}$, but also that $\tilde{\mathbf{L}} \mathbf{x} = 0$ and therefore that $\mathbf{x} \propto \mathbf{1}_n$, which implies that $\mathbf{x} = 0$. \square

Let $\tilde{\mathbf{B}} = \mathbf{B}/\mu_{\max}$, with eigenvalues

$$0 \leq 1 - \mu_i/\mu_{\max} \leq 1 - \mu < 1,$$

where $\mu = \frac{\mu_{\min}}{\mu_{\max}}$. Using Neumann expansion, $\tilde{\mathbf{B}}^{-1}$ writes:

$$\tilde{\mathbf{B}}^{-1} = \sum_{r=0}^{\infty} (\mathbf{I} - \tilde{\mathbf{B}})^r \Rightarrow \mathbf{Y}(\mathbf{A}) = \sum_{r=0}^{\infty} (\mathbf{I} - \tilde{\mathbf{B}})^r \mu_{\max}^{-1} \tilde{\mathbf{S}}_{in} \mathbf{Y}_{obs}. \quad (13)$$

We denote by \mathbf{T}_r the r -th term in $\mathbf{Y}(\mathbf{A})$:

$$\mathbf{T}_r = (\mathbf{I} - \tilde{\mathbf{B}})^r \mu_{\max}^{-1} \tilde{\mathbf{S}}_{in} \mathbf{Y}_{obs}. \quad (14)$$

Note that since $\|\tilde{\mathbf{S}}_{in} \mathbf{Y}_{obs}\| \leq \sqrt{|V_{tr}|} \|\mathbf{Y}_{obs}\|_{\infty}$, we have:

$$\|\mathbf{T}_r\| \leq \frac{\nu^r y_{\infty}}{\mu_{\max} \sqrt{|V_{tr}|}}, \quad (15)$$

where $y_{\infty} = \|\mathbf{Y}_{obs}\|_{\infty}$ and $\nu = 1 - \mu$. Similarly, $\|\mathbf{Y}(\mathbf{A})\| \leq \frac{y_{\infty}}{\mu_{\min} \sqrt{|V_{tr}|}}$. Moreover, since $\mathbf{I} - \tilde{\mathbf{B}}$ has the same zero-pattern than \mathbf{A} (except on the diagonal), if u is more than r hops from V_{tr} , we get $(\mathbf{T}_r)_u = 0$.

B. Gradient of $(\mathbf{T}_r)_u$

In the second step, we derive the formula of the gradient of $(\mathbf{T}_r)_u$ w.r.t. \mathbf{A} , and derive a bound on its magnitude as a function of r, q the distance to V_{tr} , and k the distance to V_{out} . For $r > 0$, the gradient of the u -th coefficient in \mathbf{T}_r w.r.t. $\mathbf{I} - \tilde{\mathbf{B}}$ is:

$$\begin{aligned} \nabla_{\mathbf{I} - \tilde{\mathbf{B}}} (\mathbf{T}_r)_u &= \sum_{h=1}^r \left(((\mathbf{I} - \tilde{\mathbf{B}})^{r-h})_{u,:} \right)^{\top} \\ &\quad \times \left((\mathbf{I} - \tilde{\mathbf{B}})^{h-1} \mu_{\max}^{-1} \tilde{\mathbf{S}}_{in} \mathbf{Y}_{obs} \right)^{\top}, \end{aligned}$$

by the product rule of differentiation, and we have

$$\nabla_{\mathbf{I} - \tilde{\mathbf{B}}} (\mathbf{T}_r)_u = \sum_{h=1}^r \left(((\mathbf{I} - \tilde{\mathbf{B}})^{r-h})_{u,:} \right)^{\top} (\mathbf{T}_{h-1})^{\top}.$$

Using that $\mathbf{I} - \tilde{\mathbf{B}} = \mathbf{I} - \frac{1}{\mu_{\max}} (\tilde{\mathbf{S}}_{in} + \lambda \tilde{\mathbf{L}})$, we have

$$\begin{aligned} \nabla_{\tilde{\mathbf{L}}} (\mathbf{T}_r)_u &= -\frac{\lambda}{\mu_{\max}} \nabla_{\mathbf{I} - \tilde{\mathbf{B}}} (\mathbf{T}_r)_u \\ &= -\frac{\lambda}{\mu_{\max}} \sum_{h=1}^r \left(((\mathbf{I} - \tilde{\mathbf{B}})^{r-h})_{u,:} \right)^{\top} (\mathbf{T}_{h-1})^{\top}. \end{aligned} \quad (16)$$

And finally, by deriving $\tilde{\mathbf{L}}$ w.r.t. \mathbf{A}_{ij} :

$$\begin{aligned} \frac{\partial (\mathbf{T}_r)_u}{\partial \mathbf{A}_{ij}} &= -\frac{\lambda}{|E| \mu_{\max}} \sum_{h=1}^r \left((\mathbf{I} - \tilde{\mathbf{B}})^{r-h} \right)_{ui} (\mathbf{T}_{h-1})_i \\ &\quad + \left((\mathbf{I} - \tilde{\mathbf{B}})^{r-h} \right)_{uj} (\mathbf{T}_{h-1})_j \\ &\quad - \left((\mathbf{I} - \tilde{\mathbf{B}})^{r-h} \right)_{uj} (\mathbf{T}_{h-1})_i \\ &\quad - \left((\mathbf{I} - \tilde{\mathbf{B}})^{r-h} \right)_{ui} (\mathbf{T}_{h-1})_j, \end{aligned} \quad (17)$$

which allows us to prove the following.

Lemma IV.3. *Let i, j, u such that: i, j are at least k -hop from u , and at least q -hop from V_{tr} . Then:*

$$\left| \frac{\partial (\mathbf{T}_r)_u}{\partial \mathbf{A}_{ij}} \right| \leq \begin{cases} 0 & \text{if } q + k > r \\ \frac{4\lambda y_{\infty}}{|E| \mu_{\max}^2 \sqrt{|V_{tr}|}} (r - q - k) \nu^{r-1} & \text{otherwise.} \end{cases} \quad (18)$$

Proof. Recall that $(\mathbf{T}_r)_u = 0$ if u is more than r -hop from V_{tr} . Similarly, $((\mathbf{I} - \tilde{\mathbf{B}})^r)_{ui} = 0$ if u and i are more than r -hop from each other. Hence, the term $((\mathbf{I} - \tilde{\mathbf{B}})^{r-h})_{ui} (\mathbf{T}_{h-1})_i$ appearing in (17) is 0 if $r - h < k$ or $h - 1 < q$, and bounded by $(\mu_{\max} \sqrt{|V_{tr}|})^{-1} \nu^{r-1} y_{\infty}$ otherwise. Similarly for the other terms, so the sum in (17) runs over the indices h that satisfy $q + 1 \leq h \leq r - k$, which is either none if $q + 1 + k > r$, or $r - q - k$ terms otherwise, which concludes the proof. \square

C. Proof of Theorem IV.2

We finally examine the hypergradient, and prove an exponential damping rate of its magnitude with the cumulative distance to V_{tr} and V_{out} (the sum of both distances). Considering $F_{out} = \|\mathbf{S}_{out}(\mathbf{Y}(\mathbf{A}) - \mathbf{Y}_{obs})\|^2$, where \mathbf{S}_{out} is the diagonal selection matrix whose diagonal entries equal 1 if the corresponding node is in V_{out} and 0 otherwise, we have:

$$\begin{aligned} \frac{\partial F_{out}}{\partial \mathbf{A}_{ij}} &= 2 \left(\frac{\partial \mathbf{Y}(\mathbf{A})}{\partial \mathbf{A}_{ij}} \right)^\top \mathbf{S}_{out}(\mathbf{Y}(\mathbf{A}) - \mathbf{Y}_{obs}) \\ &= 2 \sum_{r=0}^{\infty} \left(\frac{\partial \mathbf{T}_r}{\partial \mathbf{A}_{ij}} \right)^\top \mathbf{S}_{out}(\mathbf{Y}(\mathbf{A}) - \mathbf{Y}_{obs}) . \end{aligned}$$

Using a triangular inequality, the bound on $\|\mathbf{Y}(\mathbf{A})\|$, and that $\|\mathbf{S}_{out} \mathbf{Y}_{obs}\| \leq \sqrt{|V_{out}|} y_\infty$ we get:

$$\|\mathbf{S}_{out}(\mathbf{Y}(\mathbf{A}) - \mathbf{Y}_{obs})\| \leq \frac{1 + \mu_{\min} \sqrt{|V_{tr}| |V_{out}|}}{\mu_{\min} \sqrt{|V_{tr}|}} y_\infty .$$

By incorporating the resulting inequality in bounding the hypergradient, and by noticing that $\mathbf{S}_{out} = \mathbf{S}_{out}^2$ we have:

$$\begin{aligned} \left| \frac{\partial F_{out}}{\partial \mathbf{A}_{ij}} \right| &\lesssim \frac{1 + \mu_{\min} \sqrt{|V_{tr}| |V_{out}|}}{\mu_{\min} \sqrt{|V_{tr}|}} y_\infty \sum_{r=0}^{\infty} \|\mathbf{S}_{out} \frac{\partial \mathbf{T}_r}{\partial \mathbf{A}_{ij}}\| \\ &\lesssim \frac{1 + \mu_{\min} \sqrt{|V_{tr}| |V_{out}|}}{\mu_{\min} \sqrt{|V_{tr}|}} y_\infty \sum_{r=0}^{\infty} \left(\sum_{u \in V_{out}} \left| \frac{\partial (\mathbf{T}_r)_u}{\partial \mathbf{A}_{ij}} \right|^2 \right)^{\frac{1}{2}} . \end{aligned}$$

Using Lemma IV.3 and the hypotheses on i and j , for u in V_{out} , the term $\left| \frac{\partial (\mathbf{T}_r)_u}{\partial \mathbf{A}_{ij}} \right|$ is 0 if $r < q + k + 1$, and bounded by $\frac{4\lambda y_\infty}{|E| \mu_{\max}^2 \sqrt{|V_{tr}|}} (r - q - k) \nu^{r-1}$ otherwise. Hence:

$$\begin{aligned} \left| \frac{\partial F_{out}}{\partial \mathbf{A}_{ij}} \right| &\lesssim \lambda \frac{\sqrt{|V_{out}|} + \mu_{\min} \sqrt{|V_{tr}| |V_{out}|}}{\mu_{\min} |V_{tr}| |E| \mu_{\max}^2} y_\infty^2 \\ &\quad \times \sum_{r=q+k+1}^{\infty} (r - q - k) \nu^{r-1} . \end{aligned}$$

Then we see that for $\nu < 1$ we have

$$\sum_{r=q+k+1}^{\infty} (r - q - k) \nu^{r-1} = \nu^{q+k} \sum_{r=1}^{\infty} r \nu^{r-1} ,$$

and $\sum_{r=1}^{\infty} r \nu^{r-1} = \frac{1}{(1-\nu)^2} = \frac{1}{\mu^2}$, which concludes the proof.

V. ALLEVIATING HYPERGRADIENT SCARCITY

In this section, we review strategies to mitigate the hypergradient scarcity problem. However, it is important that we make a distinction between resolving this issue and resolving the overfitting problem. Indeed, if gradient scarcity is also caused by the limited quantity of available labelled data, it is important to avoid confusion with traditional overfitting. In particular, while traditional overfitting is generally reduced by adding more training data, *when optimizing edges far from labelled nodes gradient scarcity is observed regardless of the dataset size and the number of labels*. We study several strategies to mitigate hypergradient scarcity in the bilevel setting, but we emphasize that they might not lead to a better generalization error altogether.

Generalized edge refinement by optimizing \mathbf{A}_{obs}^r . As hypergradient scarcity is observed on edges connecting nodes distant from the labelled ones, a natural fix is to reduce this distance. One way to do that is by refining edges in a power of \mathbf{A}_{obs} , as the matrix \mathbf{A}_{obs}^r includes r -edge long connections between nodes. In our experiments we adopt \mathbf{A}_{obs}^6 as this notably expands the graph but does not achieve the extreme case where the result is a complete graph.

Graph regularization. Graph regularization is used to impose a prior structure on the learned graph, by adding a regularization term to F_{out} to penalize graphs with undesirable properties. For instance, [25] proposes the regularization term $-\gamma \mathbf{1}_n^\top \log \mathbf{A} \mathbf{1}_n$ for some $\gamma > 0$, to penalize low-degree nodes. We use this choice in the experiments, but note that imposing task-related priors and regularization terms could lead to better performance. This will be the topic of future work.

G2G for edge refinement. The third fix we suggest is latent graph learning using G2G models. In the outer problem, we propose to replace optimizing edge weights by optimizing the parameters of a G2G model to predict similarity between nodes. Let θ be the weights of this model, and \mathbf{A}_θ be its output graph, the G2G model we adopt is $(\mathbf{A}_\theta)_{i,j} = \alpha((\mathbf{X}_i - \mathbf{X}_j)^2)$, where the square function is applied entrywise, $\alpha: \mathbb{R}^p \rightarrow \mathbb{R}$ is a Multi-Layer Perceptron (MLP) model consisting of k_{G2G} layers, each is of the form:

$$\mathbf{X}^{[l]} = \phi^{[l]}(\mathbf{X}^{[l-1]} \mathbf{W}_1^{[l]} + \mathbf{1}_n (\mathbf{b}^{[l]})^\top) ,$$

where $\mathbf{W}_1^{[l]} \in \mathbb{R}^{d_{l-1} \times d_l}$, $\mathbf{b}^{[l]} \in \mathbb{R}^{d_l}$ are learnable parameters, and d_l is the output dimensionality of the l -th layer. The parameters are gathered as $\theta = \{\mathbf{W}_1^{[l]}, \mathbf{b}^{[l]}\}_{l=1}^{k_{G2G}}$.

VI. EXPERIMENTS

We¹ use two synthetic datasets, the first one, called synthetic dataset 1, is designed to test the Laplacian regularization. The second one is a binary classification dataset that can be used for both graph-based models. Due to the paradigm behind construction, we call it the cheaters dataset. We also illustrate our findings on the real-world Cora dataset.

Bilevel optimization: The problem Eq. (5) is intractable as neither the solution of the inner problem nor its gradient *w.r.t.* \mathbf{A} (or to θ with G2G models) has a closed form expression that can be evaluated. To overcome this difficulty, we unroll [26] τ_{in} iterations of the gradient-based inner optimizer, then using the Higher package [27] to trace iterations and perform higher-order Automatic Differentiation (AD) to compute the hypergradient. For both the inner and the outer optimizers, we consider the Adam algorithm [28].

Synthetic dataset 1: we sample *i.i.d.* latent variables $\mathbf{X} \in \mathbb{R}^{n \times p}$ for nodes uniformly at random from $[0, 1]$ with $n = 1536, p = 2$. The ground-truth graph \mathbf{A}^* is constructed s.t. $(\mathbf{A}^*)_{i,j} = 1$ if $\|\mathbf{X}_i - \mathbf{X}_j\|_2 < \sigma$, and 0 otherwise. σ is set to 0.06 in our experiments. Two distinct procedures were employed to sample the nodes that comprise V_{tr} , leading to two distinct realizations of the dataset as illustrated in Fig. 3(top). The first procedure randomly samples 100 nodes

¹Our Python implementation is available at https://github.com/hashemghanem/Gradients_scarcity_graph_learning.




Fig. 1: Hypergradient scarcity observed when solving the edge refinement task with the bilevel optimization framework. We run the experiment on the cheaters dataset, and use a 2-layer GNN as a classifier. Left: graph initialization. Right: hypergradient at an arbitrary outer iteration, namely 9. It is clear that the hypergradient on edges between unlabelled nodes far from the ones in $V_{out} \cup V_{tr}$ equals zero. Recall that $V_{tr} = \{0, 1, \dots, 32\} \cup \{224, \dots, 255\}$ and $V_{out} = \{96, \dots, 160\}$.

from the set V , hence V_{tr} is well-spread, whereas the second procedure selects the 100 nodes with the smallest Euclidean distance to the point $(0.5, 0.5)$, thus V_{tr} is concentrated in a small neighborhood in this case. In both cases, we randomly sample 25 nodes from V to construct V_{out} . The remaining nodes are equally divided between the validation and the test sets. Then, each node i in V_{tr} is labeled as follows:

$$(\mathbf{Y}_{obs})_i = \zeta \left(e^{-\frac{(\mathbf{x}_i - \mathbf{a}_1)^2}{2(0.2)^2}} + e^{-\frac{(\mathbf{x}_i - \mathbf{a}_2)^2}{2(0.2)^2}} + e^{-\frac{(\mathbf{x}_i - \mathbf{a}_3)^2}{2(0.2)^2}} \right),$$

where $\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3$ are randomly sampled from $[0, 1]^2$, and ζ is a scaling factor such that labels lie in $[0, 1]$. By this construction, the prior that the labelling function on the graph is smooth is met, and the Laplacian regularization can be applied as in Eq. (1).

To generate labels for other nodes, we plug the labels of V_{tr} and \mathbf{A}^* in Eq. (1) with $\lambda = 1$, such that the solution holds the sought-for labels. This way, the ground-truth graph actually plays a role in labelling nodes in V_{out} and V .

The noisy observed graph is built upon random weights

$$(\mathbf{A}_{obs})_{i,j} = \xi_{i,j} (\mathbf{A}^*)_{i,j} \quad \text{where} \quad \xi_{i,j} \sim \mathcal{U}([0, 1]).$$

Experiments on this dataset are done with the Laplacian regularization in the inner problem as in Eq. (1).

Cheaters dataset: nodes in this graph represent students in an exam classroom. Setting $n = 256, p = 10$, the *i.i.d.* features $\mathbf{X} \in \mathbb{R}^{256 \times 10}$ are sampled uniformly at random from $[0, 1]$. For a node i , $\mathbf{X}_{i,0}$ represents the position of the according student in the classroom. For visualization purposes we enumerate nodes following the ascending order of $\mathbf{X}_{:,0}$. The remaining 9 features of a student represent the grades he is capable of scoring in the corresponding exam question. However, students tend to cheat with their neighbors in the graph. The ground-truth graph \mathbf{A}^* is constructed as follows:

$$(\mathbf{A}^*)_{i,j} = \exp(-\|\mathbf{X}_{i,0} - \mathbf{X}_{j,0}\|_2^2 / 2\sigma^2).$$

The observed graph \mathbf{A}_{obs} is drawn from a random model as

$$(\mathbf{A}_{obs})_{i,j} \sim \text{Ber}((\mathbf{A}^*)_{i,j}).$$

We set $\sigma = 0.027$ s.t. the number of edges in \mathbf{A}_{obs} approximates $n \log n$. Students cheat such that their grades \mathbf{Y}_{grade} after the exam are

$$\mathbf{Y}_{grade} = \mathbf{A}^* \mathbf{X}_{:,1:9} \mathbf{1}_9.$$

A student passes the exam if his grade is greater than a threshold τ , *i.e.*, $(\mathbf{Y}_{obs})_i = 1$ if $(\mathbf{Y}_{grade})_i > \tau$ and 0 otherwise. We put $\tau = 60$ so that approximately half of students pass the exam. V_{tr} includes nodes in $\{0, 1, \dots, n/8\} \cup \{7n/8, \dots, n-1\}$, *i.e.*, near the two ends of the 1-dimensional class. $V_{out} = \{3n/8, \dots, 5n/8\}$, *i.e.*, centered around the middle of the class. Remaining nodes are equally divided into a validation and a test set. Experiments on this dataset are done with a GNN classifier.

Real-world dataset: we validate our findings on the Cora dataset [29]. Cora is a citation datasets, where nodes represent research publications described by a bag of words, and edges stand for citations. The task is to classify articles *w.r.t.* their topic. In this work, we limit our experimentation on real-world datasets to the Cora dataset, as our empirical results are intended to establish a proof-of-concept. Therefore, we refrain from conducting experiments on other benchmark datasets.

Models: G2G and GNN models are implemented using *PyTorch* [30] and *PyTorch Geometric* [31], respectively. The function α in the G2G model is an MLP with 2 hidden layers, each is followed by the *ReLU* activation function and has 16 neurons for the cheaters dataset and 32 neurons for Cora. The GNN has 1 hidden layer of 8 neurons for the cheaters dataset and 128 for Cora. This layer is followed *ReLU*, while the output is followed by the *softmax* function.

Setup: we use Adam as the inner and the outer optimizer with the default parameters of *PyTorch*, except for the inner learning rate η_{in} and the outer one η_{out} , which are tuned from the set $\{10^{-4}, 10^{-3}, \dots, 10\}$. The best values were $\eta_{in} = 10^{-2}$ with GNNs as a classifier, $\eta_{in} = 10^{-1}$ and $\eta_{in} = 10$ with the Laplacian regularization on Cora and on the synthetic dataset 1, respectively. On the cheaters dataset, $\eta_{out} = 10^{-3}$ adopting a G2G model, while $\eta_{out} = 10^{-2}$ in other cases. On the synthetic dataset 1, $\eta_{out} = 10^{-1}$. On Cora, $\eta_{out} = 10^{-2}$ in all experiments without a G2G model, otherwise $\eta_{out} = 10^{-4}$ adopting the GNN classifier, and $\eta_{out} = 10^{-3}$ adopting the Laplacian regularization. We set, with a grid search, τ_{in} to 200 for the cheaters dataset, 500 for the synthetic dataset 1 and Cora adopting the Laplacian regularization, and 100 for Cora with a GNN classifier. In experiments on the cheaters dataset, we multiply the default initialization of the last layer of the G2G model by 10^{-5} s.t. its output edges at the first iteration are of small magnitude. We adopt this strategy to measure the level of scarcity by counting the number of learned edges of magnitude greater than a chosen threshold. GNN weights W and the initialization of labels when using the Laplacian regularization are initialized at random after each outer iteration, using Xavier initialization and uniformly at random from $[0, 1]$, respectively. Edges to be refined are initialized uniformly at random from $[0, 1]$, except for experiments on the cheaters dataset where the interval becomes $10^{-5} * [0, 1]$. We set the number of outer iterations τ_{out} to 150 while ensuring convergence, and

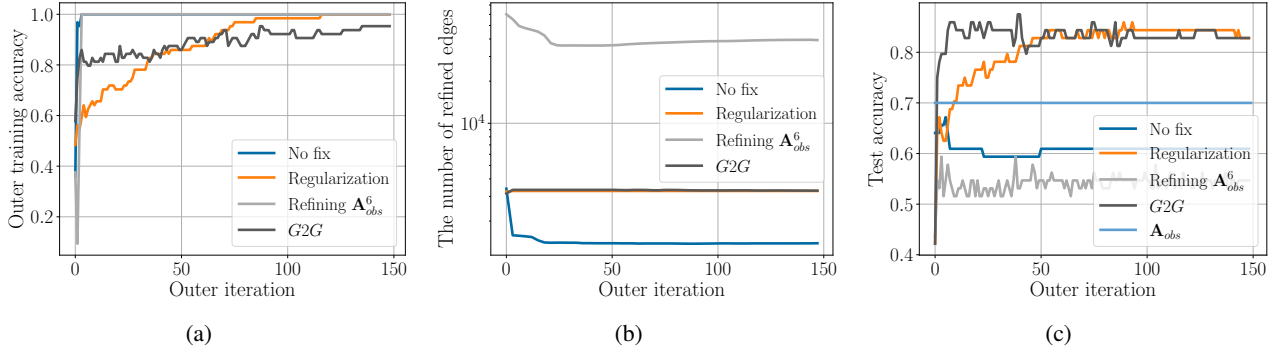


Fig. 2: Efficiency of proposed solutions to hypergradient scarcity *w.r.t.* the number of refined edges and the generalization capacity. An edge is considered well refined if its learned magnitude is larger than one percent of the maximum learned edge weight. The solutions are graph regularization with $-\mathbf{1}^\top \log \mathbf{A}\mathbf{1}$, latent graph learning using a G2G model, and generalized edge refinement by refining edges in \mathbf{A}_{obs}^6 . (a): training accuracy on V_{out} . (b): number of refined edges. (c) test accuracy.




Fig. 3: hypergradient scarcity under the bilevel optimization setting on the synthetic dataset 1, adopting the Laplacian regularization in the inner problem. **Top:** illustration of the graph. The training nodes V_{tr} are circled in red, the colors correspond to the distance to V_{tr} . The eigenvalue μ_{min} is given as a ratio of the smallest positive eigenvalue of $\tilde{\mathbf{L}}$. V_{out} is randomly sampled from V but not shown here. **Bottom:** Hypergradient magnitude $\left| \frac{\partial F_{out}}{\partial \mathbf{A}_{ij}} \right|$ with respect to the *sum* of distances to V_{tr} and V_{out} . **Left:** the training set V_{tr} is well-spread thereby aligned with the high-frequency eigenvectors of the graph, resulting in a *high* μ_{min} . The decrease of the hypergradients is sharp with the distance. **Right:** V_{tr} is aligned with the low-frequency eigenvectors of the graph, resulting in a *low* μ_{min} . The decrease of hypergradients magnitude is not as sharp as the previous case.

we select the graph (or the G2G weights) with the highest validation accuracy. We set $\lambda = 1$ in training when considering the Laplacian regularization, as we expect the bilevel algorithm to learn this parameter by scaling the learned adjacency matrix. When applying the Laplacian regularization fed with \mathbf{A}_{obs} on

Cora, we set $\lambda = 0.1$ after a grid search. γ in the graph regularization term is set to 1 following a grid search on the set $\{10^{-3}, 10^{-2}, \dots, 10\}$.

A. Hypergradient scarcity with GNN classifiers

In this experiment, we consider a 2-layer GNN classifier in the bilevel framework. We solve the edge refinement task (5) on the cheaters dataset, where ℓ in Eqs. (1) and (5) is the CCE function. Fig. 1(left) depicts the initialization of the adjacency matrix. It also shows what edges are to be optimized, that is, edges whose initialization is nonzero. In Fig. 1(right), we show the hypergradient at the outer iteration 9, which is arbitrarily chosen, where it is clear that edges between unlabelled nodes far from the ones in the union $V_{out} \cup V_{tr}$ get no supervision during the training process. Recall that $V_{tr} = \{0, 1, \dots, 32\} \cup \{224, \dots, 255\}$ and $V_{out} = \{96, \dots, 160\}$. This aligns with our findings, which state that edges between nodes at least 2-hop from nodes in $V_{out} \cup V_{tr}$ receive zero hypergradients. This, as seen in Fig. 2, leads to a learned graph that overfits training nodes and even generalizes worse than \mathbf{A}_{obs} .

B. Hypergradient scarcity with the Laplacian regularization

We here examine hypergradient scarcity when adopting the Laplacian regularization in the inner problem. We run the bilevel optimizer to solve the edge refinement task on the synthetic dataset 1. The dataset corresponds to a regression problem, so ℓ in Eqs. (1) and (5) is the MSE loss function.

In Fig. 3(bottom), we plot the absolute value of hypergradients at the outer iteration 6 as a function of the edge cumulative distance to V_{tr} and V_{out} , which is defined as follows: we compute $q + k$, the sum of distances to V_{tr} and V_{out} , respectively, for its both endpoint nodes, then we take the minimum of the two results. One observes the hypergradient scarcity phenomenon, since hypergradients decay exponentially as the edge distance increases. This validates our analysis articulated in Theorem IV.2. In addition, we observe in practice that μ is nevertheless quite small, and that our bound in Theorem IV.2 is quite loose. Another observation is that the decrease rate is higher when V_{tr} is well-spread in the graph. Deriving a tighter

TABLE I: Accuracies obtained on Cora when the classifier is trained using the output graph of the Bilevel Optimization (BO) framework, the same framework equipped with graph regularization, the same framework optimizing a G2G model. We also benchmark against GAM (the result is reported from the according paper) and against \mathcal{A}_{obs} . For each classifier, we report test accuracy in the according first line and training accuracy on V_{out} in the second one. Training accuracy on V_{tr} equals 100% for all methods.

Graph	\mathcal{A}_{obs}	BO	BO+regularization	BO+G2G	GAM
GNN	77.0	76.2	80.3	82.0	84.8
	77.4	94.9	94.1	97.4	-
Laplacian	71.7	76.2	78.3	76.2	-
	71.0	81.9	83.2	83.5	-

bound on the magnitude of hypergradients and investigating the link between the distribution of labelled nodes and this bound will be the subject of a future work.

C. Testing solutions to mitigate hypergradient scarcity

We run our experiments on the cheaters dataset using the 2-layer GNN as a classifier. In each experiment, we run our bilevel optimization framework with one of the suggested fixes. We consider two criteria to measure the efficiency of each solution, the first one is counting the number of refined edges. At any outer iteration, we say that an edge is refined if its learned weight is greater than one percent of the maximum learned edge weight at the same iteration. Recall that we initialize the graph/GtoG with small weights ($\approx 10^{-5}$). The second criterion is the validation accuracy. The first criterion assesses the ability to alleviate hypergradient scarcity, while the second assesses the generalization to unseen nodes during training, and thus if the learned graph is meaningful.

Fig. 2 shows that all three fixes produce better results *w.r.t.* the first criterion, as the number of refined edges is larger at almost every iteration, with optimizing edges in \mathcal{A}_{obs}^6 being the most efficient, and the G2G model and graph regularization having a similar performance. Moreover, one notices that this number decreases with the iteration when refining edges (without fix) in \mathcal{A}_{obs}^6 , which is expected as only a small portion of edges receive supervision; however this portion is larger when refining \mathcal{A}_{obs}^6 .

Regarding the second criterion, the G2G model and the graph regularization generalize well, as both combat hypergradient scarcity without increasing (or even by decreasing) the number of parameters to learn. On the other hand, optimizing edges in \mathcal{A}_{obs}^6 deteriorates performance in the test phase. A likely explanation is that by expanding the graph, we increase the number of parameters to learn, which means a more complex model that is more likely to overfit training nodes. This experiment illustrates that **hypergradient scarcity is not the traditional overfitting** related to data/label scarcity, and resolving it does not necessarily promote better generalization.




Fig. 4: Observing hypergradient scarcity and the effect of graph regularization on Cora. Left: adopting the GNN classifier. Right: adopting the Laplacian regularization model. We plot the hypergradient against edge distance. In connected components without at least a node from each of V_{tr} and V_{out} in the Laplacian regularization case (or without a node from $V_{tr} \cup V_{out}$ in the GNN case), edge distance is not defined. We assign the distance 15 to edges in such components for visualization purpose.

D. Results on Cora

We use bilevel optimization (5) to solve an edge refinement task on Cora, trying both the GNN and the Laplacian models. Here the downstream task is a multi-label classification problem and ℓ is the CCE function. We depict in Fig. 4 the received hypergradient on edges at outer iteration 9 as a function of their distance to labelled nodes. For the Laplacian regularization case, that is the edge cumulative distance to V_{tr} and V_{out} as defined in Section VI-B. Whereas to compute the edge distance in the GNN case, we compute for each of its endpoint nodes its distance to $V_{tr} \cup V_{out}$, then we take the minimum. In accordance with our analysis, the figure displays a null hypergradient for distances greater than 2 in the GNN case, while the Laplacian regularization scenario exhibits a hypergradient that diminishes exponentially with distance.

Regarding the generalization capacity, Table I shows that the learned graph is inferior to \mathcal{A}_{obs} in the GNN case for the test error. Given that the learned graph achieves 100%, 94.9% accuracies on V_{tr} , V_{out} , respectively, one concludes that hypergradient scarcity provokes overfitting. This is indeed expected due to the extreme scarcity in the GNN scenario, as edges of distance greater than 2 keep their random initialization after the training process. This is, however, not the case in the Laplacian regularization scenario as most edges are of distance less than 11, thereby they do not exhibit damped hypergradients and the impact on generalization is not observed.

Next, we test the efficiency of the proposed solutions to mitigate hypergradient scarcity. We do not try learning a power of \mathcal{A}_{obs} as the memory requirement goes beyond the limits we have access to. Results in Fig. 4 prove the efficiency of graph regularization as all edges receive non-zero hypergradients with a comparable magnitude to those on edges of small distance. Note that hypergradients are received on the G2G weights when it is deployed, not on edges, so we do not depict them in this figure. Regarding the impact on generalization, Table I shows that both fixes yield significant improvements in test accuracy over \mathcal{A}_{obs} with the GNN classifier. In the Laplacian regularization case, graph regularization produces a

higher test accuracy, unlike the G2G model which generalizes equally good as when learning directly edge weights. We also notice that GNN model leads to superior results in all scenarios, with a notable gap for the G2G model and graph regularization, and when directly using \mathbf{A}_{obs} . This is expected, as the Laplacian regularization promotes similarity between connected nodes but, unlike GNNs, is not a supervised-based method. We finally point out that the bilevel optimization framework with either fix does not achieve state-of-the-art results produced by GAM with the same GNN classifier.

Other experiments suggest that although G2G models alleviate hypergradient scarcity, regardless of the number of neurons in its layers, the generalization performance is sensitive to this number and if set large, clear overfitting is observed.

VII. CONCLUSION

We studied hypergradient scarcity when deploying bilevel optimization in edge refinement tasks under the SSL settings. This phenomenon consists in edges far from labelled nodes receiving scarce hypergradients when optimizing the graph and the classifier to improve classification performance. We proved that this problem occurs for GNN. Replacing GNNs by the Laplacian regularization model, does not resolve the issue; however, the phenomenon is less severe: we bounded the magnitude of hypergradients and proved they are exponentially damped with distance to labelled nodes. To alleviate hypergradient scarcity, we proposed to resort to latent graph learning, graph regularization, and refining edges in a power of the observed adjacency matrix. Our experiments validated our findings, and privileged the first two solutions over the latter. Moreover, we show that alleviating the hypergradient scarcity does not necessarily alleviate overfitting.

REFERENCES

- [1] F. Wang and C. Zhang, "Label propagation through linear neighborhoods," in *ICML*, 2006.
- [2] Q. Liu, M. Allamanis, M. Brockschmidt, and A. Gaunt, "Constrained graph variational autoencoders for molecule design," *Advances in neural information processing systems*, vol. 31, 2018.
- [3] D. Liben-Nowell and J. Kleinberg, "The link prediction problem for social networks," in *Proceedings of the twelfth international conference on Information and knowledge management*, 2003, pp. 556–559.
- [4] B. Fatemi, L. El Asri, and S. M. Kazemi, "Slaps: Self-supervision improves structure learning for graph neural networks," *NeurIPS*, 2021.
- [5] K. P. Bennett, J. Hu, X. Ji, G. Kunapuli, and J.-S. Pang, "Model selection via bilevel optimization," in *IJCNN*, 2006.
- [6] R. Flamary, A. Rakotomamonjy, and G. Gasso, "Learning constrained task similarities in graph regularized multi-task learning," in *Regularization, Optimization, Kernels, and Support Vector Machines*. Chapman and Hall/CRC, 2014, vol. 103.
- [7] L. Franceschi, P. Frasconi, S. Salzo, R. Grazzi, and M. Pontil, "Bilevel programming for hyperparameter optimization and meta-learning," in *ICML*, 2018.
- [8] B. Colson, P. Marcotte, and G. Savard, "An overview of bilevel optimization," *Ann. Oper. Res.*, vol. 153, no. 1, pp. 235–256, 2007.
- [9] S. T. Roweis and L. K. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *Science*, vol. 290, no. 5500, pp. 2323–2326, 2000.
- [10] J. B. Tenenbaum, V. d. Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, no. 5500, pp. 2319–2323, 2000.
- [11] L. Franceschi, M. Niepert, M. Pontil, and X. He, "Learning discrete structures for graph neural networks," in *ICML*, 2019, pp. 1972–1982.
- [12] O. Stretcu, K. Viswanathan, D. Movshovitz-Attias, E. Platanios, S. Ravi, and A. Tomkins, "Graph agreement models for semi-supervised learning," *NeurIPS*, 2019.
- [13] T. Luong, H. Pham, and C. D. Manning, "Effective approaches to attention-based neural machine translation," in *EMNLP*, 2015.
- [14] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *NeurIPS*, 2017.
- [15] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph Attention Networks," *ICLR*, 2018.
- [16] D. Kim and A. Oh, "How to find your friendly neighborhood: Graph attention design with self-supervision," in *ICLR*, 2021.
- [17] H. Wang and J. Leskovec, "Unifying graph convolutional neural networks and label propagation," *arXiv preprint arXiv:2002.06755*, 2020.
- [18] X. Zhu, *Semi-supervised learning with graphs*. Carnegie Mellon University, 2005.
- [19] N. Keriven, "Not too little, not too much: a theoretical analysis of graph (over) smoothing," in *NeurIPS*, 2022.
- [20] X. Liu, F. Zhang, Z. Hou, L. Mian, Z. Wang, J. Zhang, and J. Tang, "Self-supervised learning: Generative or contrastive," *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- [21] L. Wu, H. Lin, C. Tan, Z. Gao, and S. Z. Li, "Self-supervised learning on graphs: Contrastive, generative, or predictive," *IEEE Trans. Knowl. Data Eng.*, 2021.
- [22] Y. Liu, M. Jin, S. Pan, C. Zhou, Y. Zheng, F. Xia, and P. Yu, "Graph self-supervised learning: A survey," *IEEE Trans. Knowl. Data Eng.*, 2022.
- [23] Y. Liu, Y. Zheng, D. Zhang, H. Chen, H. Peng, and S. Pan, "Towards unsupervised deep graph structure learning," in *WWW*, 2022.
- [24] G. W. Stewart, *Matrix algorithms: volume 1: basic decompositions*. SIAM, 1998.
- [25] V. Kalofolias, "How to learn a graph from smooth signals," in *Artificial Intelligence and Statistics*. PMLR, 2016, pp. 920–929.
- [26] K. Gregor and Y. LeCun, "Learning fast approximations of sparse coding," in *Proceedings of the 27th international conference on international conference on machine learning*, 2010, pp. 399–406.
- [27] E. Grefenstette, B. Amos, D. Yarats, P. M. Htut, A. Molchanov, F. Meier, D. Kiela, K. Cho, and S. Chintala, "Generalized inner loop meta-learning," *arXiv preprint arXiv:1910.01727*, 2019.
- [28] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *International Conference on Learning Representations*, 12 2014.
- [29] Q. Lu and L. Getoor, "Link-based classification," in *ICML*, 2003.
- [30] A. Paszke and al., "Pytorch: An imperative style, high-performance deep learning library," in *NeurIPS*, 2019.
- [31] M. Fey and J. E. Lenssen, "Fast graph representation learning with PyTorch Geometric," in *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.