

RESEARCH

Open Access



# Enhancing molecular property prediction with auxiliary learning and task-specific adaptation

Vishal Dey<sup>1</sup> and Xia Ning<sup>1,2,3\*</sup>

## Abstract

Pretrained Graph Neural Networks have been widely adopted for various molecular property prediction tasks. Despite their ability to encode structural and relational features of molecules, traditional fine-tuning of such pretrained GNNs on the target task can lead to poor generalization. To address this, we explore the adaptation of pretrained GNNs to the target task by jointly training them with multiple auxiliary tasks. This could enable the GNNs to learn both general and task-specific features, which may benefit the target task. However, a major challenge is to determine the relatedness of auxiliary tasks with the target task. To address this, we investigate multiple strategies to measure the relevance of auxiliary tasks and integrate such tasks by adaptively combining task gradients or by learning task weights via bi-level optimization. Additionally, we propose a novel gradient surgery-based approach, Rotation of Conflicting Gradients (RCG<sub>rad</sub>), that learns to align conflicting auxiliary task gradients through rotation. Our experiments with state-of-the-art pretrained GNNs demonstrate the efficacy of our proposed methods, with improvements of up to 7.7% over fine-tuning. This suggests that incorporating auxiliary tasks along with target task fine-tuning can be an effective way to improve the generalizability of pretrained GNNs for molecular property prediction.

## Scientific contribution

We introduce a novel framework for adapting pretrained GNNs to molecular tasks using auxiliary learning to address the critical issue of negative transfer. Leveraging novel gradient surgery techniques such as RCG<sub>rad</sub>, the proposed adaptation framework represents a significant departure from the dominant pretraining fine-tuning approach for molecular GNNs. Our contributions are significant for drug discovery research, especially for tasks with limited data, filling a notable gap in the efficient adaptation of pretrained models for molecular GNNs.

**Keywords** Pretraining, Graph neural networks, Auxiliary learning, Task adaptation, Molecular property prediction, Drug discovery

## Introduction

Accurate prediction of molecular properties is pivotal in drug discovery [39], as it accelerates the identification of potential molecules with desired properties. Developing computational models for property prediction relies on learning effective representations of molecules [5]. In this regard, Graph Neural Networks (GNNs) have shown impressive results in learning effective representations for molecular property prediction tasks [11, 12, 37]. Inspired by the paradigm of pretraining followed

\*Correspondence:

Xia Ning

ning.104@osu.edu

<sup>1</sup> Department of Computer Science and Engineering, The Ohio State University, Columbus 43210, OH, USA

<sup>2</sup> Department of Biomedical Informatics, The Ohio State University, Columbus 43210, OH, USA

<sup>3</sup> Translational Data Analytics Institute, The Ohio State University, Columbus 43210, OH, USA



© The Author(s) 2024. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.

by fine-tuning, widely recognized for its impact in natural language understanding [27, 38], molecular GNNs are often pretrained [17] on a large corpus of molecules. Such a corpus might encompass irrelevant data for the target property prediction task. This can lead the GNNs to learn features that do not benefit the target task. Consequently, pretrained GNNs are fine-tuned with the target task to encode task-specific features. However, vanilla fine-tuning can potentially lead to poor generalization, particularly when dealing with diverse downstream tasks, limited data, and the need to generalize across varying scaffolds [40].

To improve generalization, auxiliary learning has recently garnered attention [8, 20, 21]. Auxiliary learning leverages informative signals from self-supervised tasks on unlabeled data, to improve the performance of the target tasks. However, its application in the context of molecular graphs, specifically for molecular property prediction, remains largely unexplored. Following this line of work, in this paper, we explore how to adapt pretrained molecular GNNs by combining widely-used self-supervised tasks with the target task using respective task-specific data (with self-supervised and target task labels). However, a critical challenge in such an adaptation is caused by negative transfer [29], where auxiliary tasks might impede rather than aid the target task [9, 30].

To address this challenge, we develop novel gradient surgery-based adaptation strategies, referred to as Rotation of Conflicting Gradients (RCGrad) and Bi-level Optimization with Gradient Rotation (BLO+RCGrad). Such strategies mitigate negative transfer from auxiliary tasks by learning to align conflicting gradients. Overall, our adaptation strategies improved the target task performance by as much as 7.7% over vanilla fine-tuning. Moreover, our findings indicate that the developed adaptation strategies are particularly effective in tasks with limited labeled data, which is a common challenge in molecular property prediction tasks. Our comprehensive investigation of multiple adaptation strategies for pretrained molecular GNNs represents a notable contribution in addressing the limited benefit of pretrained GNNs [34], and in improving generalizability across a diverse set of downstream tasks with limited data.

## Related work

### Pretraining and fine-tuning GNNs

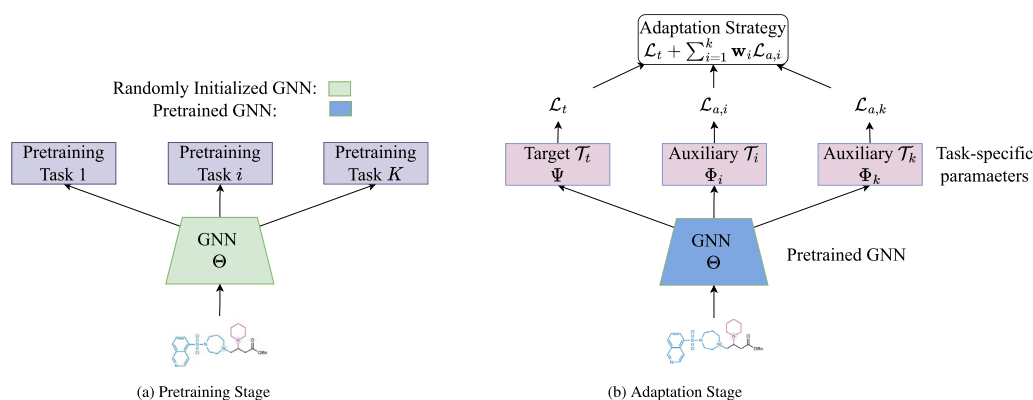
Pretraining followed by fine-tuning is widely used to leverage knowledge gained from related tasks and to improve model generalization. Typically, it involves training a model on large-scale data with self-supervised or supervised tasks, and then fine-tuning it on a small-scale labeled data. Following the success of pretraining

and fine-tuning paradigm in various domains [10, 23], researchers have extended it to molecular GNNs [17, 18, 22, 37]. In this regard, researchers have designed a number of self-supervised tasks as pretraining tasks that focus on capturing diverse chemical rules, connectivities, and patterns at varying granularities: on node, subgraph and graph levels [42]. Although pretrained GNNs showed promise in capturing diverse chemical knowledge, the challenge lies in effectively extracting this knowledge relevant to the target task, which is often non-trivial through vanilla fine-tuning. Specifically, such fine-tuning often leads to overfitting [41]. Contrary to the observations in domains such as natural language processing (NLP) and computer vision, where pretrained models consistently yield substantial improvements, pretrained GNNs do not exhibit such improvement [34].

This could be due to a notable research gap in determining what self-supervised molecular tasks can better benefit the downstream target tasks. In fact, prior studies in pretraining molecular GNNs mostly leverage one or two self-supervised task(s), thereby resulting in a plethora of multiple pretrained GNNs. Interestingly, such pretrained GNNs capture different knowledge [36] and excel in different downstream molecular property prediction tasks [34]. Additionally, Sun et al. [34] recently demonstrated that self-supervised graph pretraining does not consistently/significantly outperform non-pretraining methods across various settings. Overall, although pretrained GNNs hold promise for molecular property prediction, their benefit over non-pretrained models seems limited. To address this, some recent attempts [41, 46] to fine-tune pretrained GNNs have largely relied on existing ideas like regularization [43] or update constraints [16] during fine-tuning. In contrast, our proposed approaches leverage auxiliary tasks to learn generalizable knowledge and prevent overfitting to the training set.

### Knowledge transfer with auxiliary learning

Knowledge transfer through auxiliary learning has demonstrated its effectiveness across a spectrum of domains [19, 26, 35]. This paradigm, distinct from multi-task learning, aims to optimize the target task's performance while leveraging auxiliary tasks to bolster generalization [32]. Prior research in other domains has developed multiple methods to automatically learn task weights, such as using gradient similarity [6, 9], using parameterized auxiliary network [8, 25], using bi-level optimization and implicit differentiation [2, 25], minimizing distances between task embeddings [3], or from the perspective of Nash equilibrium [31]. However, the application of auxiliary learning for adapting molecular GNNs to target tasks, particularly in the context of



**Fig. 1** Off-the-shelf available pretrained GNNs are transferred for target task-specific adaptation

molecular property prediction, remains an under-explored area. In this study, we adopt and explore gradient similarity, gradient scaling, and bi-level optimization strategies.

## Preliminaries

Motivated by the success of continued pretraining and task-specific adaptation in pretrained Large Language Models (LLMs) [7, 13, 44], we investigate adaptation of off-the-shelf pretrained molecular GNNs to target molecular property prediction tasks. Via such an adaptation, we aim to leverage existing self-supervised (SSL) tasks designed for molecular GNNs and transfer learned knowledge from such tasks to the target task. We employ the existing SSL tasks typically used in molecular pretraining such as masked atom prediction (AM), context prediction (CP) [17], edge prediction (EP) [14], graph infomax (IG) [33], and motif prediction (MP) [28]. (detailed in B.1). We refer to these tasks as auxiliary tasks. Intuitively, these auxiliary tasks can potentially capture diverse chemical semantics and rich structural patterns at varying granularities. By utilizing SSL objectives on target task-specific data, auxiliary tasks augment the pretrained GNNs with richer representations. Such representations, in turn, can improve the generalizability of the target property prediction task. Henceforth, the term “GNN” refers to an off-the-shelf pretrained molecular GNN.

Figure 1 presents an overview of the adaptation setup. Formally, we adapt a GNN with parameters  $\Theta$  to optimize the performance on the target task  $\mathcal{T}_t$ . We achieve this by jointly training  $\mathcal{T}_t$  with auxiliary tasks  $\{\mathcal{T}_{a,i}\}_{i=1}^k$  through solving the following optimization problem:

$$\min_{\Theta, \Psi, \Phi_{i \in \{1..k\}}} \mathcal{L}_t + \sum_{i=1}^k \mathbf{w}_i \mathcal{L}_{a,i}, \quad (1)$$

where  $\mathcal{L}_t$  and  $\mathcal{L}_{a,i}$  denote the target task loss and  $i$ -th auxiliary task loss, respectively,  $\Psi$  and  $\Phi_{i \in \{1, \dots, k\}}$  denotes

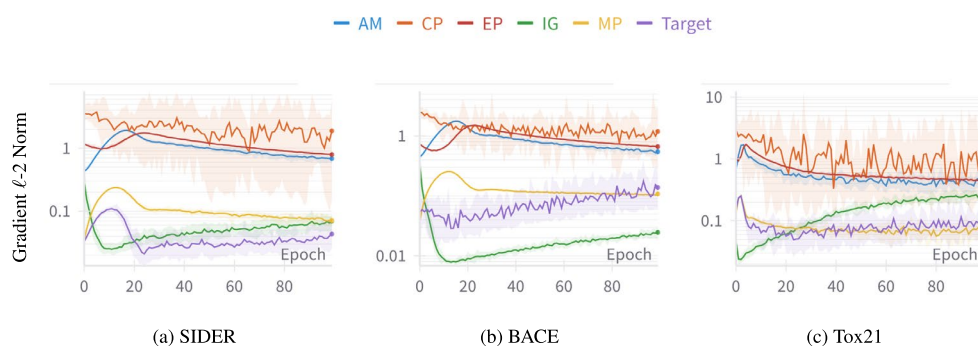
task-specific learnable parameters for the target and  $i$ -th auxiliary task, respectively, and  $\mathbf{w}$  is the weight indicating the influence of the auxiliary tasks on the target task. Through the above optimization, all the parameters are simultaneously updated in an end-to-end manner. Note that the above optimization does not optimize  $\mathbf{w}$ — we will introduce an approach that can additionally learn  $\mathbf{w}$  in Section “Bi-Level Optimization”. In fact, the key to effective adaptation lies in accurately determining  $\mathbf{w}$ , such that the combined task gradients can backpropagate relevant training signals to the shared GNN as follows:

$$\Theta^{(t+1)} := \Theta^{(t)} - \alpha \left( \mathbf{g}_t + \sum_{i=1}^k \mathbf{w}_i \mathbf{g}_{a,i} \right),$$

where  $\mathbf{g}_t = \nabla_{\Theta} \mathcal{L}_t$ , and  $\mathbf{g}_{a,i} = \nabla_{\Theta} \mathcal{L}_{a,i}$  denote the gradients updating  $\Theta$  from the target and  $i$ -th auxiliary task, respectively, and  $\alpha$  denotes the learning rate. Our proposed adaptation strategies focus on learning such  $\mathbf{w}$  in an end-to-end manner, to dynamically combine task gradients during each update. These strategies contrast with those using fixed weights or conducting expensive grid-search to explore all possible  $\mathbf{w}$ .

## Gradient cosine similarity (GCS)

The first strategy to meaningfully combine task gradients is based on gradient cosine similarity (GCS) [9]. Intuitively, GCS measures the alignment between task gradients during training, providing insights into the relatedness of auxiliary tasks with the target task. A high GCS indicates that the auxiliary tasks provide complementary information, and thus, can benefit the target task. Conversely, low GCS indicates potential orthogonality or even conflict between tasks. Thus, GCS can naturally quantify the relatedness of auxiliary tasks with the target task over the course of training. We compute GCS and update  $\Theta$  as:



**Fig. 2** Large variations of scales among task gradients are observed when Sup-CP is adapted with all auxiliary tasks using MTL

$$\Theta^{(t+1)} := \Theta^{(t)} - \alpha \left( \mathbf{g}_t + \sum_{i=1}^k \max(0, \cos(\mathbf{g}_t, \mathbf{g}_{a,i})) \mathbf{g}_{a,i} \right),$$

where, max operator takes the maximum out of the two values, thereby, dropping the tasks with conflicting gradients (i.e., with negative GCS).

#### Gradient scaling (GNS)

We also adopt a simpler strategy of gradient scaling [15] to adjust the influence of auxiliary tasks with respect to the target task. Our preliminary experiments as presented in Figure 2 revealed significant differences in the scales of the task gradient norms, and thus requiring careful adjustments. This is because if the gradient of an auxiliary task is much larger than that of the target task,  $\Theta$  updates will be most dominated by such auxiliary tasks, thereby potentially resulting in worse target performance. On the other hand, if the gradient of an auxiliary task is relatively small, the training signals from such auxiliary tasks will be too weak to encode any relevant features in  $\Theta$ . Thus, following [4, 15], we use a simple gradient scaling to dynamically adjust the influence of auxiliary tasks during updates of  $\Theta$  as follows:

$$\Theta^{(t+1)} := \Theta^{(t)} - \alpha \left( \mathbf{g}_t + \sum_{i=1}^k \max \left( 1, \frac{\|\mathbf{g}_t\|}{\|\mathbf{g}_{a,i}\|} \right) \mathbf{g}_{a,i} \right), \quad (2)$$

where  $\|\cdot\|$  denotes the  $\ell$ -2 norm.

## Methods

#### Rotation of conflicting gradients (RCGrad)


While both conflicting directions and magnitude differences of task gradients can lead to negative transfer, GCS and GNS focus separately on homogenizing either the direction or magnitude of gradients, rather than in a unified manner. To address these limitations, we develop Rotation of Conflicting Gradients (RCGrad)

– a novel extension of PCGrad [45] – that aligns gradients both in terms of direction and magnitude. RCGrad, which builds upon PCGrad, does not completely discard gradients conflicting with the target task, unlike GCS. Instead, RCGrad only negates the component of the conflicting gradient that is completely opposite to the target task gradient. Additionally, RCGrad explicitly learns how much of the non-conflicting component should be incorporated for the most effective knowledge transfer. This mitigates negative transfer by not only removing the conflicting component but also by learning to incorporate a portion of the non-conflicting component.

Figure 3 demonstrates the difference between PCGrad and RCGrad. Formally, RCGrad learns to rotate auxiliary gradient  $\mathbf{g}_{a,i}$  by angle  $\theta_i$  to yield a rotated gradient  $\mathbf{R}(\theta_i) \mathbf{g}_{a,i}$ , which is followed by an orthogonal projection in case of conflicts (Figure 3b). The orthogonally projected component is computed as  $\mathbf{g}_{a,i}^r = \text{oproj}_t \mathbf{R}(\theta_i) \mathbf{g}_{a,i}$ , where  $\mathbf{R}(\theta_i)$  is the rotation matrix parameterized by  $\theta_i$ , and  $\text{oproj}_t$  is the orthogonal vector projection operator as defined in Eq. 3. Via such an operator (Figure 3a), PCGrad projects the conflicting auxiliary gradient  $\mathbf{g}_{a,i}$  onto the normal plane of the target task's gradient  $\mathbf{g}_t$  to yield  $\mathbf{g}_{a,i}^p$  as follows:

$$\mathbf{g}_{a,i}^p = \text{oproj}_t \mathbf{g}_{a,i} = \mathbf{g}_{a,i} - \frac{\mathbf{g}_{a,i} \cdot \mathbf{g}_t}{\|\mathbf{g}_t\|} \cdot \frac{\mathbf{g}_t}{\|\mathbf{g}_t\|}, \quad (3)$$

where  $\text{oproj}_t$  denotes the orthogonal projection operator with respect to  $\mathbf{g}_t$ . This enables effective knowledge transfer from auxiliary tasks, even if they share some dissimilarity to the target task. However, PCGrad does not explicitly learn how much of the non-conflicting component should be incorporated for the most effective knowledge transfer. To address this limitation, RCGrad learns an appropriate rotation to be applied to the auxiliary gradient  $\mathbf{g}_{a,i}$ , followed by the projection of the rotated gradient. Such a learnable rotation in an end-to-end manner



**Fig. 3** **a** PCGrad projects conflicting gradient  $\mathbf{g}_{a,i}$  onto the normal plane of  $\mathbf{g}_t$ . **b** RCGrad applies a rotation to  $\mathbf{g}_{a,i}$ , followed by projection. **c** Rotation followed by orthogonal projection is equivalent to scaling  $\mathbf{g}_{a,i}^p$ . **d** If the rotated gradient does not conflict with  $\mathbf{g}_t$ , the projection of the rotated gradient onto  $\mathbf{g}_t$  is incorporated as scaling  $\mathbf{g}_t$  by  $(1 + \mathbf{s}_t)$

enables dynamic knowledge transfer from auxiliary tasks such that the target task performance can be improved.

Moreover, as shown in Figures 3c, d, the rotation followed by the projection of gradients is equivalent to applying appropriate scaling factors  $\mathbf{s}_i$  and  $\mathbf{s}_t$  on the projected gradients  $\mathbf{g}_{a,i}^p$  and  $\mathbf{g}_t$ , respectively. Additionally, different from PCGrad, RCGrad accounts for large differences in gradient magnitudes by adjusting the magnitudes of non-conflicting auxiliary task gradients relative to that of the target task gradient (Eq. 2). To summarize,  $\Theta$  is updated as follows:  $\Theta^{(t+1)} := \Theta^{(t)} - \alpha \mathbf{g}$ , where  $\mathbf{g} = (1 + \mathbf{s}_t) \times \mathbf{g}_t + \sum_{i=1}^k \mathbf{g}_{a,i}^r$ , and

$$\mathbf{g}_{a,i}^r = \begin{cases} \mathbf{s}_i \times \mathbf{g}_{a,i}^p & \text{if } \mathbf{g}_t \cdot \mathbf{g}_{a,i} < 0, \\ \max\left(1, \frac{\|\mathbf{g}_t\|}{\|\mathbf{g}_{a,i}\|}\right) \mathbf{g}_{a,i} & \text{otherwise,} \end{cases} \quad (4)$$

where  $\mathbf{g}_{a,i}^p$  is computed via Eq. 3. Note that the set of scaling factors  $\mathbf{s} = \{\{\mathbf{s}_i\}_{i=1}^k, \mathbf{s}_t\}$  is learned in an end-to-end manner during the optimization of the combined losses from all tasks.

### Bi-level optimization (BLO)

Unlike the previous approaches that directly manipulate task gradients, BLO learns task weights  $\mathbf{w}$  (Eq. 1) in an end-to-end manner, such that the GNN generalizes well to the target task. Note that BLO does not directly intervene in the gradient computation process. Instead, BLO learns  $\mathbf{w}$  that minimizes the target validation loss while ensuring that the GNN is optimized with a weighted combination of losses:

$$\mathbf{w}^* = \arg \min_{\mathbf{w}} \mathcal{L}_t^{(A)}(\Theta^*(\mathbf{w})), \quad \text{s.t.} \quad \Theta^*(\mathbf{w}) = \arg \min_{\Theta} \mathcal{L}_f(\Theta, \mathbf{w}) \quad (5)$$

where,  $\mathcal{L}_f = \mathcal{L}_t + \sum_{i=1}^k \mathbf{w}_i \mathcal{L}_{a,i}$  is the combined loss on the training set, and  $\mathcal{L}_t^{(A)}$  is the loss on the target task computed with a held-out auxiliary dataset  $\mathcal{A}$ , and  $\Theta^*(\mathbf{w})$  is the best-response of  $\Theta$  with current  $\mathbf{w}$ . This formulation is a bi-level optimization problem: updating  $\mathbf{w}$  in the upper-level optimization requires

computing  $\nabla_{\mathbf{w}} \mathcal{L}_t^{(A)} = \nabla_{\Theta} \mathcal{L}_t^{(A)} \cdot \nabla_{\mathbf{w}} \Theta^*$ , where the latter gradient requires back-propagation through the inner-level optimization of  $\Theta$ . Following [24], we leverage the Implicit Function Theorem (IFT) to compute  $\nabla_{\mathbf{w}} \Theta^* = -(\nabla_{\Theta}^2 \mathcal{L}_f)^{-1} \cdot \nabla_{\mathbf{w}} \nabla_{\Theta} \mathcal{L}_f$ . Intuitively, IFT allows us to evaluate the  $\nabla_{\mathbf{w}} \Theta^*$  locally around the approximate best-response  $\Theta^*$ . Using the above, we can compute the gradients  $\nabla_{\mathbf{w}} \mathcal{L}_t^{(A)}$  as:

$$\begin{aligned} \nabla_{\mathbf{w}} \mathcal{L}_t^{(A)}(\Theta^*(\mathbf{w})) &= \nabla_{\Theta} \mathcal{L}_t^{(A)} \cdot \nabla_{\mathbf{w}} \Theta^*(\mathbf{w}) \\ &= -\nabla_{\Theta} \mathcal{L}_t^{(A)} \cdot (\nabla_{\Theta}^2 \mathcal{L}_f)^{-1} \cdot \nabla_{\mathbf{w}} \nabla_{\Theta} \mathcal{L}_f. \end{aligned} \quad (6)$$

We described the entire training process in Algorithm 1 (Appendix A). To compute the Hessian inverse and vector products efficiently, we use the iterative algorithm by Lorraine et al. [24], which is summarized in Algorithm 2 (Appendix A). Intuitively, it uses a Neumann series expansion to approximate the Hessian inverse with unrolling differentiation for  $M$  steps around locally approximate best-response  $\Theta^*$ . Following [25], in practice, we don't train  $\Theta$  till convergence (i.e.,  $\Theta^*$  such that  $\nabla_{\Theta} \mathcal{L}_f = 0$ ). Instead, we approximate  $\Theta^*$  by simultaneously training both  $\Theta$  and  $\mathbf{w}$ , and alternately optimizing  $\mathbf{w}$  for every  $r$  updates of  $\Theta$ . We refer the readers to [24] for theoretical considerations on approximations and convergence. Note that we use 20% of the training set as  $\mathcal{A}$  instead of using the validation set to avoid data leakage and unfair comparison with baselines. Optimizing  $\mathbf{w}$  on a held-out  $\mathcal{A}$  rather than on

the training set aligns with the goal of improving target task generalizability.

### BLO with gradient rotation (BLO+RCGrad)

In the previous sections, we discussed RCGrad, which learns to project and scale conflicting gradients using

**Table 1** Test ROC-AUC using  $\mathcal{T}_a=\{\text{AM,CP,EP,IG,MP}\}$  and Sup-CP

Method	SIDER	ClinTox	BACE	BBBP	Tox21	ToxCast	HIV	MUV
FT	61.82 (0.53)	<b>71.10</b> (1.40)	82.86 (0.87)	67.57 (1.39)	77.05 (0.34)	66.02 (0.18)	78.70 (0.80)	80.64(0.51)
GTOT	62.24 (0.34)	70.03 (1.58)	83.67 (1.75)	69.01 (1.95)	<b>77.08</b> (0.66)	65.45 (0.45)	<b>80.05</b> (0.57)	82.09(3.13)
MTL	56.22 (0.82)	56.41 (3.43)	80.04 (1.48)	64.88(1.23)	74.42 (0.34)	64.53 (0.38)	76.79 (0.26)	81.68 (0.49)
GCS	59.94 (0.53)	62.77 (2.17)	85.60 (0.63)	71.16(0.52)	74.76 (0.42)	66.05 (0.19)	76.94 (1.30)	76.65 (1.17)
GNS	62.48 (0.56)	67.94 (1.02)	84.80 (0.34)	70.94(0.86)	76.44 (0.24)	66.19 (0.21)	78.23 (0.44)	<b>83.50</b> (1.14)
PCGrad	62.09 (0.62)	67.60 (1.88)	84.42 (1.23)	69.14 (1.25)	76.58 (0.77)	65.81 (0.61)	77.76 (1.08)	78.62 (0.29)
BLO	60.70 (2.37)	68.29 (3.02)	85.14 (1.23) <sup>a</sup>	69.80 (0.68)	76.57 (0.44)	65.80 (0.79)	79.16 (0.47)	82.19(1.40)
RCGrad	<b>62.49</b> (0.65)	<b>70.07</b> (1.70) <sup>b</sup>	<b>85.65</b> (0.60) <sup>a</sup>	<b>72.35</b> (1.28) <sup>ab</sup>	<b>77.26</b> (0.38) <sup>b</sup>	<b>66.49</b> (0.30) <sup>ab</sup>	<b>79.39</b> (0.63)	83.07 (1.26)
BLO+RCGrad	<b>62.94</b> (0.66) <sup>a</sup>	69.59 (2.12) <sup>b</sup>	<b>86.10</b> (0.35) <sup>a</sup>	<b>71.81</b> (1.57) <sup>a</sup>	76.62 (0.29)	<b>66.38</b> (0.29) <sup>a</sup>	79.03 (1.12)	<b>83.92</b> (1.03)

We report the mean (and standard deviation) over 10 different seeds with scaffold splitting. Best- and second best-performing models are in **bold** and bold. Tasks are presented in increasing order of size. a and b indicate statistical significance compared to the best finetuning and GS baselines, respectively. Statistical significance is determined based on the Wilcoxon signed rank test with  $p < 0.05$

s, and BLO, which learns task weights  $\mathbf{w}$  but does not explicitly handle gradient conflicts. In this section, we introduce a novel approach BLO+RCGrad that combines the strengths of both RCGrad and BLO. Instead of learning the scaling factors  $\mathbf{s}$  by minimizing the combined loss on the training split as in RCGrad, BLO+RCGrad learns  $\mathbf{s}$  that minimizes the target validation loss, which is similar to the optimization of  $\mathbf{w}$  in BLO. This enables learning  $\mathbf{s}$  that can effectively homogenize conflicting task gradients based on the generalization performance of the target task. In BLO+RCGrad, the bi-level optimization is employed for learning  $\mathbf{s}$  not to balance task losses but to best align conflicting task gradients. This addresses the limitation of BLO in handling gradient conflicts by incorporating the rotational alignment strategy of RCGrad. To summarize, BLO+RCGrad leverages the learned scaling factors  $\mathbf{s}$  via BLO (Algorithm 1) to guide the gradient surgery process introduced by RCGrad (Eq. 4). This dynamically controls the knowledge transfer from auxiliary tasks, ensuring that the influence of each task is optimally tuned to benefit the target task learning.

## Results and discussion

### Experimental setup


We perform experiments on 8 benchmark classification datasets from MoleculeNet [40]. We compare our adaptation strategies with simple baselines such as traditional fine-tuning (FT), and vanilla multi-task learning (MTL) that assigns equal weights to all auxiliary tasks; and a more advanced state-of-the-art regularization-based fine-tuning with optimal transport (GTOT) [46]. Additionally, we consider other state-of-the-art gradient surgery-based methods (GCS, GNS, PCGrad) as baselines. We refer to this group of baselines collectively as

GS methods. We use the official publicly available checkpoints<sup>1</sup> of two GNNs: 1) supervised\_contextpred [17], denoted as Sup-CP, which is pretrained via self-supervised context prediction and supervised graph-level multi-task learning, and 2) supervised [17], denoted as Sup, which is pretrained only via supervised graph-level multi-task learning. Using such different pretrained GNNs allows a controlled comparison to understand how different pretraining objectives (with and without self-supervised context prediction task) can influence the adaptation. Details on auxiliary tasks and datasets are presented in Table 4 and Appendix B.

### Reproducibility and implementation details

Following the prior line of research [17, 22], we use scaffold-split for the downstream target tasks, and use the same atom and bond features as in GTOT. All experimental details for the FT baseline follow the GTOT fine-tuning setup. Specifically, we initialized a linear projection layer on top of the pretrained GNN as the target task classifier. Across all methods, both the pretrained GNN and task-specific layers are trainable. For FT and adaptation methods, we train the models for 100 epochs with Adam optimizer with an initial learning rate  $\alpha$  of 0.001, we use a batch size of {32, 64, 256}, an embedding dimension of 300, and a dropout probability of 0.5 for the GNN module. For GTOT experiments, we use the optimal hyperparameters provided for each dataset, when finetuned on Sup-CP. For MTL experiments, we assign equal weights to all auxiliary tasks. For BLO and BLO+RCGrad experiments, we use  $M = 3$  in Algorithm 2, update  $\mathbf{w}$  every  $r = \{5, 10, 20\}$  update of  $\Theta$ , and use Adam optimizer with learning rate  $\beta$  of 0.001 to update  $\mathbf{w}$ . The code is available at <https://github.com/vishaldeyiest/GraphTA>.

<sup>1</sup> <https://github.com/snap-stanford/pretrain-gnns>



**Fig. 4** Target task gradient conflicts with EP and CP tasks. Sup-CPs adapted with all auxiliary tasks in a MTL setting

### Comparison using Sup-CPs as the pretrained GNN

Table 1 presents an overall comparison when all the auxiliary tasks are used with Sup-CPs as the pretrained GNN. Our proposed adaptation strategies, specifically RCGrad and BLO+RCGrad, outperform all baselines, including other GS-based adaptation strategies, across all datasets (except ClinTox). Specifically, compared to the best fine-tuning method, GTOT, RCGrad demonstrated significant improvement of 2.4% and 4.8% in BACE and BBBP, respectively. This indicates the efficacy of our proposed rotational alignment in mitigating negative transfer and improving the generalizability of the pretrained GNN. Furthermore, BLO+RCGrad exhibits significant improvement over fine-tuning methods FT and GTOT in small-scale datasets of as much as 6.3% and 4.1%, respectively. This highlights the efficacy of bi-level optimization combined with gradient rotation in improving generalizability, especially in limited data regimes.

Additionally, RCGrad and BLO+RCGrad consistently outperform other gradient surgery-based (GS) methods. Specifically, compared to PCGrad, RCGrad demonstrates statistically significant improvements in ROC-AUC by 2.5%, 4.7%, 0.9% and 1.0% in ClinTox, BBBP, Tox21, and ToxCast, respectively. This improvement can be attributed to the rotation component in RCGrad, which not only resolves gradient conflicts but also actively aligns them in a direction favorable to the target task. Moreover, our proposed methods RCGrad and BLO+RCGrad learn to retain a component of the conflicting task gradients, unlike GCS which completely discards conflicting gradients. This ensures that valuable information from auxiliary tasks is not discarded, thus facilitating more effective knowledge transfer.

Conversely, BLO, which learns task weights without explicitly handling gradient conflicts, performs comparably or slightly worse than RCGrad, BLO+RCGrad, and other GS-based baselines. The suboptimal performance of BLO, especially in smaller datasets (e.g., SIDER), may be attributed to the noisy nature of task gradients, potentially leading to a poor approximation of hyper-gradients. In contrast, GNS is more robust to

noisy gradients since it adjusts the scale of gradient magnitudes relative to the target task. Overall, our proposed methods consistently outperform all baselines on smaller datasets (except ClinTox), while achieving competitive performance on larger ones.

In contrast, MTL, which assigns equal weights to all auxiliary tasks regardless of their relevance to the target task, results in worse performance across all downstream tasks. Compared to FT, MTL exhibits deteriorations of as much as 9.1% and 20.6% in SIDER and ClinTox, respectively. This indicates that MTL leads to drastic negative transfer, where the auxiliary tasks hurt the performance of the target task. On the contrary, all adaptation strategies (including GS-based baselines) perform better than MTL with significant improvements of up to 24.2%. Furthermore, upon analyzing gradient similarities of auxiliary tasks with the target task (Figure 4), we hypothesize that AM, IG, and MP may benefit the target task better than the other auxiliary tasks.

Table 2 presents an overall comparison using only AM, IG, and MP as auxiliary tasks. Compared to fine-tuning-based methods (FT and GTOT), our proposed methods RCGrad and BLO+RCGrad demonstrate better performance across 6 out of 8 datasets. Specifically, compared to GTOT, RCGrad achieves significant improvements of 2.6% and 5.4% in BACE and BBBP, respectively. Furthermore, RCGrad and BLO+RCGrad exhibit better performance than GS baselines with significantly improved ROC-AUC of as much as 9.9% in ClinTox. Overall, our proposed methods demonstrate significantly improved performance in smaller datasets compared to fine-tuning and GS baselines. Such consistently superior performance underscores the robustness of our methods, particularly in settings where data is limited and the alignment of gradients is crucial.

In contrast with the previous setup, GS baselines such as GCS and GNS exhibit better performance across almost all datasets. This implies that these methods can be more effective with fewer conflicting tasks, and may struggle to handle a large number of conflicting tasks (Table 1). Similarly, with fewer tasks in this setup, MTL

**Table 2** Test ROC-AUC using  $\mathcal{T}_a=\{\text{AM,IG,MP}\}$  and Sup-CP

Method	SIDER	ClinTox	BACE	BBBP	Tox21	ToxCast	HIV	MUV
FT	61.82 (0.53)	<b>71.10</b> (1.40)	82.86 (0.87)	67.57 (1.39)	77.05 (0.34)	66.02 (0.18)	78.70 (0.80)	80.64 (0.51)
GTOT	62.24 (0.34)	70.03 (1.58)	83.67 (1.75)	69.01 (1.95)	<b>77.08</b> (0.66)	65.45 (0.45)	<b>80.05</b> (0.57)	82.09 (3.13)
MTL	59.15 (1.84)	62.01 (1.87)	83.60 (0.43)	71.67 (4.44)	75.64 (0.37)	65.14 (0.21)	78.18 (1.07)	81.26 (1.90)
GCS	62.83 (0.70)	64.62 (1.83)	84.17 (0.87)	70.49 (4.26)	<b>77.35</b> (0.20)	66.03 (0.12)	77.59 (1.24)	80.17 (3.26)
GNS	62.62 (0.49)	63.42 (2.19)	84.29 (0.97)	71.79 (3.72)	76.50 (0.39)	66.12 (0.20)	78.25 (0.60)	82.42 (0.47)
PCGrad	61.42 (1.69)	63.44 (2.90)	83.92 (1.23)	70.86 (4.54)	76.73 (0.89)	65.96 (0.71)	77.38 (1.10)	80.45 (2.34)
BLO	<b>62.85</b> (0.77)	67.52 (3.27) <sup>b</sup>	<b>84.79</b> (0.62)	71.93 (3.19) <sup>a</sup>	76.88 (0.26)	<b>66.29</b> (0.26) <sup>a</sup>	79.21 (0.33)	81.86 (1.19)
RCGrad	<b>63.12</b> (0.38) <sup>ab</sup>	69.91 (1.22) <sup>b</sup>	<b>85.86</b> (0.38) <sup>ab</sup>	<b>72.76</b> (1.05) <sup>aa</sup>	76.86 (0.38)	<b>66.37</b> (0.16) <sup>ab</sup>	79.17 (0.32)	<b>82.68</b> (1.92)
BLO+RCGrad	62.44 (0.28)	<b>70.99</b> (2.31) <sup>b</sup>	84.75 (0.53)	<b>72.03</b> (3.83) <sup>a</sup>	76.64 (0.28)	66.25 (0.22) <sup>ab</sup>	<b>79.75</b> (0.81)	<b>82.65</b> (3.36)

Best- and second best-performing models are in **bold** and bold. a and b indicate statistical significance compared to the best baselines based on the Wilcoxon signed rank test with  $p < 0.05$

**Table 3** Test ROC-AUC using  $\mathcal{T}_a=\{\text{AM,CPE,IG,MP}\}$  and Sup

Method	SIDER	ClinTox	BACE	BBBP	Tox21	ToxCast	HIV	MUV
FT	61.85 (0.68)	54.16 (5.25)	75.76 (0.65)	66.34 (0.82)	75.64 (0.22)	63.52 (0.23)	72.84 (0.85)	<b>80.46</b> (0.19)
GTOT	<b>62.38</b> (0.39)	55.64 (7.49)	75.82 (2.10)	66.26 (1.87)	75.25 (1.11)	64.00 (0.55)	74.93 (1.50)	<b>80.42</b> (0.42)
MTL	55.18 (0.96)	47.33 (1.84)	64.84 (2.43)	63.62 (1.08)	73.15 (0.44)	62.06 (2.00)	63.25 (5.15)	69.21 (8.51)
GCS	58.39 (0.59)	50.05 (1.48)	74.59 (0.61)	66.67 (2.41)	74.36 (0.43)	63.94 (0.35)	72.23 (0.24)	62.99 (5.35)
GNS	60.57 (2.04)	53.52 (5.44)	76.69 (0.88)	68.67 (0.42)	75.37 (0.34)	63.49 (0.12)	74.41 (0.19)	79.72 (0.17)
PCGrad	59.83 (0.53)	53.07 (5.12)	71.17 (6.65)	67.18 (1.12)	74.26 (0.53)	63.95 (0.42)	71.80 (0.45)	79.31 (0.74)
BLO	60.65 (2.66)	56.10 (4.77)	75.11 (1.19)	67.81 (1.09) <sup>a</sup>	74.57 (0.59)	<b>64.20</b> (0.44)	75.05 (0.74) <sup>b</sup>	78.12 (0.68)
RCGrad	61.38 (0.74)	<b>57.36</b> (3.75)	<b>77.00</b> (1.03)	<b>68.73</b> (0.76) <sup>a</sup>	<b>75.67</b> (0.49)	63.91 (0.23)	<b>75.60</b> (0.26) <sup>b</sup>	79.37 (1.74)
BLO+RCGrad	<b>62.41</b> (0.81)	<b>59.45</b> (3.33) <sup>b</sup>	<b>77.47</b> (0.79)	<b>69.45</b> (0.70) <sup>a</sup>	<b>76.08</b> (0.34) <sup>b</sup>	<b>64.60</b> (0.28) <sup>ab</sup>	<b>75.80</b> (0.41) <sup>b</sup>	79.97 (1.11)

Best- and second best-performing models are in **bold** and bold. a and b indicate statistical significance compared to the best baselines based on the Wilcoxon signed rank test with  $p < 0.05$

exhibits improved performance compared to the previous setup, thereby indicating diminished negative transfer. This suggests that a smaller and more focused set of auxiliary tasks can lead to more efficient and less conflicting learning dynamics. However, PCGrad, RCGrad, and BLO+RCGrad, which partially utilize conflicting gradients, show mixed responses to the reduction in the number of auxiliary tasks in this setup. Specifically, RCGrad demonstrates improved performance in smaller datasets (except ClinTox) but a slight decrease in performance in larger datasets, compared to their performance in the previous setup. This can be attributed to the reduced diversity in learning signals provided by a smaller set of auxiliary tasks.

### Comparison using Sup as the pretrained GNN

Table 3 presents an overall comparison of adaptation of Sup as the pretrained GNN using all auxiliary tasks. Similar to our findings in the previous section, MTL again results in worse performance compared to fine-tuning

methods, thus indicating negative transfer. On the other hand, our proposed methods, specifically RCGrad and BLO+RCGrad, demonstrate improved performance over fine-tuning and GS baselines. Notably, compared to the best fine-tuning baseline GTOT, BLO+RCGrad improved ROC-AUC by 6.8%, 2.2%, and 4.8% in ClinTox, BACE, and BBBP, respectively. Similarly, compared to the best GS baseline GNS, BLO+RCGrad demonstrates notable improvement of 3.0%, 11.1%, and 1.0% in SIDER, ClinTox, and BACE, respectively. Furthermore, compared to BLO, which does not explicitly handle conflicting task gradients, BLO+RCGrad yields consistent improvement across most datasets. Such consistently superior performance of BLO+RCGrad implies that aligning and extracting informative components out of conflicting task gradients is crucial to improve the generalizability of pretrained GNNs, regardless of the specific pretraining objective.

Following the similar setup of Sup-CP experiments with a selected subset of auxiliary tasks, Table 5 in Appendix C presents an overall comparison using Sup



**Table 4** Overview of benchmark molecular property prediction datasets

Dataset	BBBP	Tox21	ToxCast	SIDER	ClinTox	MUV	HIV	BACE
No. mols	2,039	7,831	8,575	1,427	1,478	93,087	41,127	1,513
No. tasks	1	12	617	27	2	17	1	1
Avg. atoms	24.06	18.57	18.78	33.64	26.16	24.23	25.51	34.09
Avg. diameter	11.32	9.62	9.49	14.14	12.39	12.79	11.98	15.22

as the pretrained GNN. Compared to the previous setup with all auxiliary tasks, almost all GS baselines and our proposed method RCGrad exhibit improved performance with fewer auxiliary tasks. This suggests that using a smaller and relevant set of auxiliary tasks can lead to more efficient adaptation, which holds true across different pretrained GNNs. Furthermore, compared to the best GS baseline, GNS, our proposed methods RCGrad and BLO+RCGrad achieve better or comparable performance, particularly on smaller datasets. Additionally, BLO+RCGrad exhibits significant improvement over GCS in Tox21 and ToxCast.

However, it's worth noting that when using Sup as the pretrained GNN, all methods, including RCGrad and BLO+RCGrad, yield slightly worse performance compared to when Sup-CP is used as the pretrained GNN. This observation suggests that the Sup pretrained GNN might not capture contextual chemical relationships as effectively as Sup-CP, which was pretrained additionally on the context prediction task. This subtle difference in performance indicates that the choice of pretrained GNN can have an impact on the overall adaptation process. Additional results are presented in Appendix C.

## Conclusion and future work

In this study, we explored multiple adaptation strategies to improve the performance of pretrained GNNs on downstream molecular property prediction tasks. To address the poor generalization performance to such diverse downstream tasks, we introduced two novel methods, RCGrad and BLO+RCGrad, that learn to align conflicting task gradients. Our experiments demonstrate that our proposed methods consistently outperform all fine-tuning and gradient surgery-based approaches, especially on smaller datasets (except ClinTox). This suggests that the adaptation of pretrained GNNs can be a promising direction to boost target task performance, especially with limited labeled data. Our

study serves as the first step in exploring the adaptation of pretrained GNNs in molecular property prediction. In future work, we will explore other adaptation strategies to alleviate noisy gradients and to improve task selection with sparser task weights. We will further investigate the benefit of adapting GNNs to diverse downstream molecular regression tasks.

## Appendix A Details on BLO

The pseudocode for BLO is provided below:

### Algorithm 1 Learning task weights with BLO

---

```

1: Input:  $N, r, \alpha$ 
2: Initialize  $\mathbf{w}$  with  $1/k$ ,  $\Theta$  from pretrained GNN,  $\Psi$  and  $\Phi$  with default Xavier initializer
3: for  $epoch$  from 1 to  $N$  do
4:   Compute  $\mathcal{L}_f = \mathcal{L}_t + \sum_{i=1}^k \mathbf{w}_i \mathcal{L}_{a,i}$ 
5:    $\Theta \leftarrow \Theta - \alpha \nabla_{\Theta} \mathcal{L}_f$ ,  $\Phi \leftarrow \Phi - \alpha \nabla_{\Phi} \mathcal{L}_a$ ,  $\Psi \leftarrow \Psi - \alpha \nabla_{\Psi} \mathcal{L}_t$ 
6:   if  $epoch \% r == 0$  then
7:      $\mathbf{w} \leftarrow \mathbf{w} - \nabla_{\mathbf{w}} \mathcal{L}_t^{(A)}(\Theta(\mathbf{w}))$  ▷ Algorithm 2
8:   end if
9: end for
10: Return  $\Theta, \mathbf{w}$ 

```

---

### Algorithm 2 Computing $\nabla_{\mathbf{w}} \mathcal{L}_t^{(A)}(\Theta(\mathbf{w}))$

---

```

1: Input:  $\mathcal{L}_f, \mathcal{L}_t^{(A)}$ , current  $\mathbf{w}, \Theta$  from Algorithm 1,  $M, \beta$ 
2: Initialize  $p = q = \nabla_{\Theta} \mathcal{L}_t^{(A)}|_{(\mathbf{w}, \Theta)}$  ▷ Hessian inverse approximation
3: for  $j$  from 1 to  $M$  do
4:    $p = p - \beta p \nabla_{\Theta}^2 \mathcal{L}_f$ 
5:    $q = q + p$ 
6: end for
7: Return  $-q \nabla_{\mathbf{w}} \nabla_{\Theta} \mathcal{L}_f|_{(\mathbf{w}, \Theta)}$ 

```

---

Notably, the pseudocode for BLO+RCGrad is very similar to that of BLO, except that the gradients  $\mathbf{g}$  for BLO (in Step 5 of Algorithm 1) is computed as  $\mathbf{g} = \nabla_{\Theta} \mathcal{L}_f$ , while  $\mathbf{g}$  for BLO+RCGrad is computed as described above.

## Appendix B experimental details

### B.1 On auxiliary tasks

We describe the auxiliary tasks and share key insights behind using them:

- Masked Atom Prediction (AM): AM [17] involves predicting the identity of masked atoms within a molecular graph. It helps the GNN to learn the local chemical context and relationships between atoms and bonds, which are crucial for understanding molecular structure and function. The embedding out of GNN is fed to a linear classifier to predict the atom type of masked atoms.
- Edge Prediction (EP): EP [14] focuses on predicting the presence or absence of bonds (edges) between pairs of atoms in a molecular graph. It helps the GNN to capture essential local structural information, including connectivity and spatial arrangement of atoms within molecules. Following existing design [34], the dot product of node embeddings is used to predict the existence of a bond.
- Context Prediction (CP): CP [17] requires the model to predict neighboring graph structures (context) based on an anchor structure. This aids the GNN in distinguishing molecular contexts, enabling the model to capture subgraph-level information. The setup of Hu et al. [17] is followed to extract and distinguish positive and negative subgraph contexts.
- Graph Infomax (IG): IG [33] maximizes the mutual information between local (node) and global (subgraph) representations. This helps the GNN to capture structural patterns, allowing it to understand how atoms form functional groups and larger molecular substructures. The existing setup [33] is followed to train a discriminator model that distinguishes between node embeddings from the same molecular graph and those from a different graph.
- Motif Prediction (MP): MP [28] focuses on predicting the presence of specific recurring substructures (motifs) within a molecule. It helps the GNN to identify structural motifs indicative of chemical properties or functions. This task is formulated as a multi-label binary classification problem with each of 85 motifs<sup>2</sup> extracted from RDKit [1] as labels.

Each of these tasks focuses on different aspects of molecular graphs, such as local connectivity, spatial arrangement, contextual information, hierarchical organization, and recurring structural patterns. In essence, these tasks are designed to equip the model with a richer understanding of molecular structures, ultimately improving its ability to generalize and make accurate predictions.

Note that designing auxiliary tasks is beyond the scope of this study.

### B.2 Dataset overview

We perform our adaptation experiments on 8 benchmark classification datasets from MoleculeNet [40]. In this section, we give a brief overview and provide preliminary statistics of these datasets.

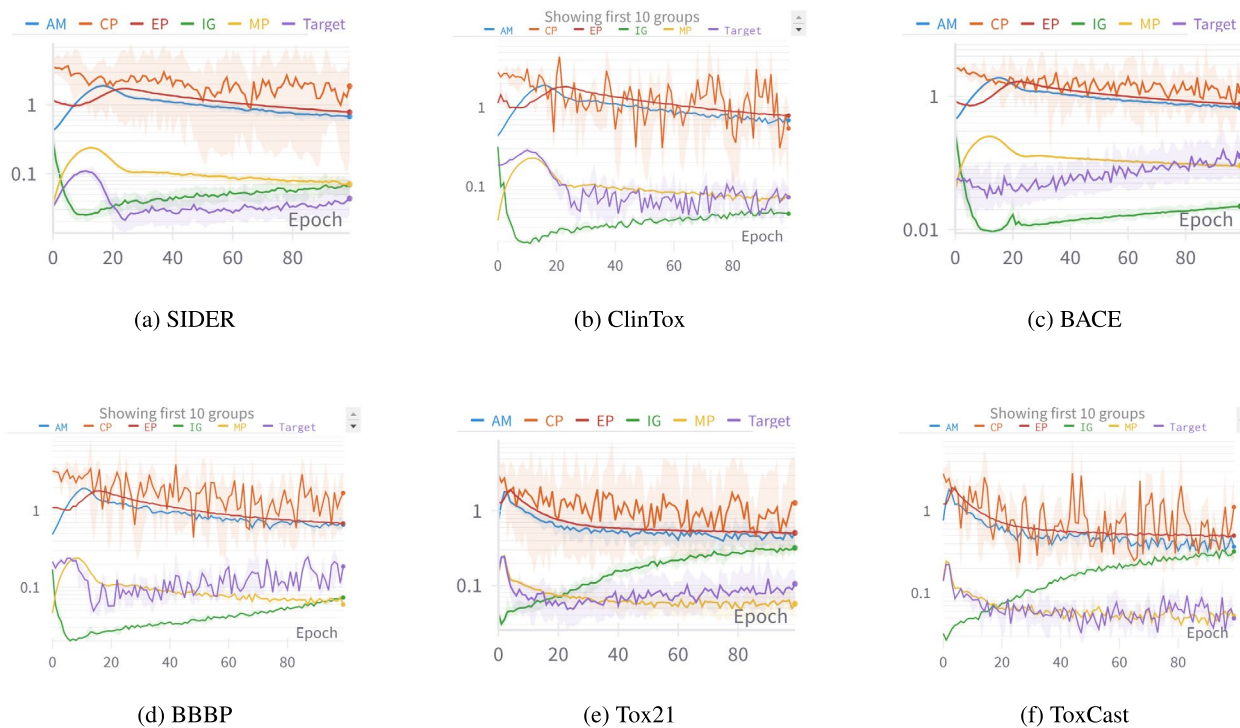
- BBBP: measures whether a molecule permeates the blood-brain barrier.
- BACE: measures whether a molecule inhibit the  $\beta$ -secretase 1 (BACE-1) enzyme.
- ClinTox: contains toxicity labels for clinical drugs, facilitating the assessment of drug safety profiles across various targets. It is important to note that these labels reflect both FDA approval outcomes and clinical trial failures due to toxicity. Such outcomes are determined by not just the molecular structures of the drugs, but also by external factors such as genetic predispositions, evaluation methodologies, and environmental conditions. This complexity can make methodological comparisons challenging.
- HIV: measures whether a molecule can prevent anti-viral activity against the HIV virus.
- MUV: compiled and refined from PubChem bioassays, evaluating compound activity across multiple targets.
- Tox21: measures toxicity across a range of biological pathways used in the 2014 Tox21 challenge.
- ToxCast: measures compound toxicity across a range of biological systems.

### B.3 Additional figures

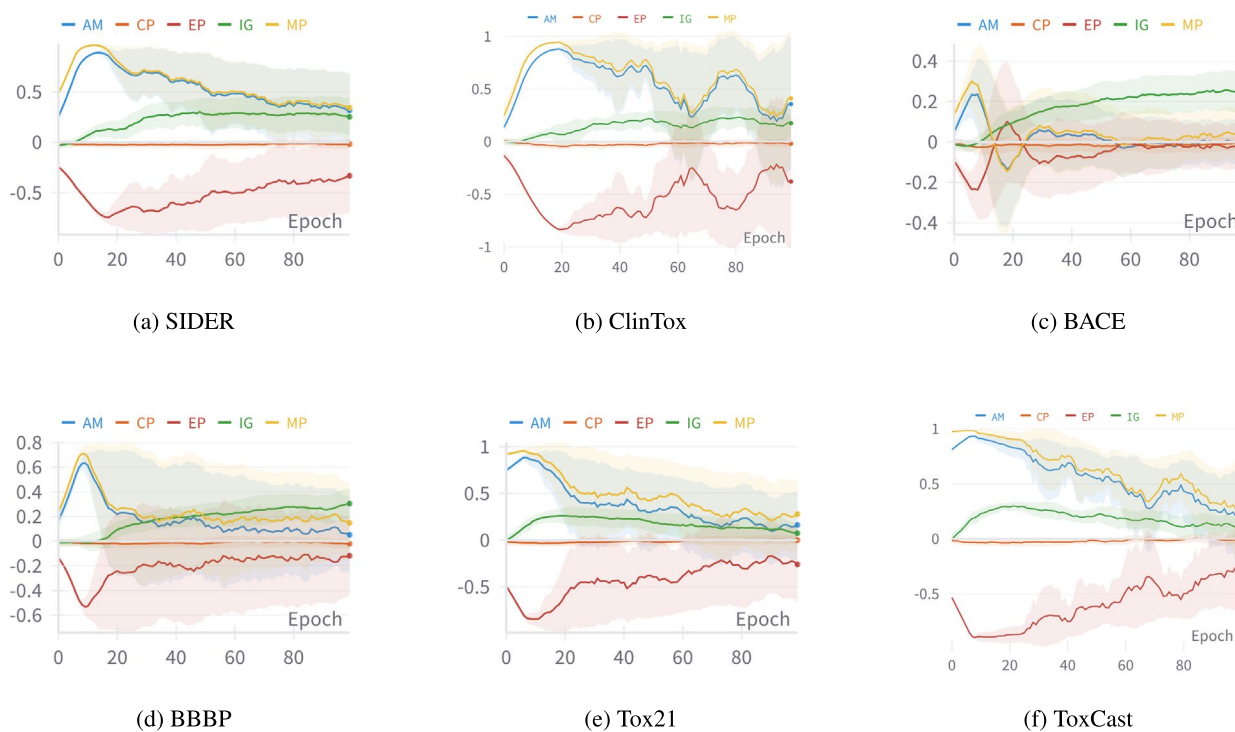
Figure 5 demonstrates the varying scales of auxiliary task gradient magnitudes when Sup-CPis adapted using all auxiliary tasks in a MTL setting across all datasets. This indicates the need to adjust the gradient norms as proposed in GNS and RCGrad. This prevents some auxiliary tasks to dominate over target tasks.

Figure 6 demonstrates that target task gradient conflicts with that of EP and CP tasks across all datasets. This motivates our experimental comparison of all adaptation strategies using a smaller set of more relevant auxiliary tasks.

<sup>2</sup> <http://rdkit.org/docs/source/rdkit.Chem.Fragments.html>



**Fig. 5** Large variations of scales among task gradients observed across multiple tasks



**Fig. 6** Target task gradient conflicts with EP and CP tasks. Sup-CPs are adapted with all auxiliary tasks in a MTL setting

**Table 5** Test ROC-AUC using  $\mathcal{T}_a=\{\text{AM,IG,MP}\}$  and Sup

Method	SIDER	ClinTox	BACE	BBBP	Tox21	ToxCast	HIV	MUV
FT	61.85 (0.68)	54.16 (5.25)	75.76 (0.65)	66.34 (0.82)	<b>75.64</b> (0.22)	63.52 (0.23)	72.84 (0.85)	<b>80.46</b> (0.19)
GTOT	62.38 (0.39)	55.64 (7.49)	75.82 (2.10)	66.26 (1.87)	75.25 (1.11)	64.00 (0.55)	74.93 (1.50)	<b>80.42</b> (0.42)
MTL	56.24 (2.79)	53.25 (2.60)	75.92 (1.06)	68.72 (0.73)	72.22 (0.62)	62.94 (0.32)	71.84 (0.95)	74.81 (0.48)
GCS	61.31 (0.65)	50.22 (1.60)	75.54 (1.12)	65.23 (1.89)	75.01 (0.30)	64.45 (0.27)	74.03 (0.52)	75.20 (1.99)
GNS	<b>62.47</b> (0.49)	55.08 (4.51)	<b>77.28</b> (1.38)	69.55 (1.10)	74.95 (0.41)	63.94 (0.23)	74.13 (0.34)	77.05 (1.53)
PCGrad	57.66 (2.41)	52.20 (3.22)	76.55 (0.96)	69.11 (0.62)	73.01 (0.90)	63.59 (0.48)	71.88 (1.15)	75.28 (1.32)
BLO	61.70 (0.86)	56.79 (3.67)	75.25 (1.54)	68.00 (0.88)*	74.53 (0.33)	64.44 (0.73)	75.15 (0.35) <sup>b</sup>	76.97 (3.12)
RCGrad	62.10 (1.04)	<b>58.64</b> (1.66) <sup>b</sup>	<b>77.64</b> (0.80)	<b>69.63</b> (0.84) <sup>a</sup>	75.08 (0.53)	<b>65.09</b> (0.37) <sup>ab</sup>	<b>75.63</b> (0.24) <sup>b</sup>	78.08 (2.78)
BLO+RCGrad	<b>62.55</b> (0.85)	<b>59.31</b> (3.59) <sup>b</sup>	77.22 (1.53)	<b>69.67</b> (0.93) <sup>a</sup>	<b>75.74</b> (0.58) <sup>b</sup>	<b>65.18</b> (0.44) <sup>ab</sup>	<b>75.78</b> (0.21) <sup>b</sup>	<b>78.37</b> (2.31)

Best- and second best-performing models are in **bold** and bold. a and b indicate statistical significance compared to the best baselines based on the Wilcoxon signed rank test with  $p < 0.05$

## Appendix C Additional experimental results

Table 5 presents an overall comparison when Sup is adapted using only AM, IG, and MP as auxiliary tasks. Compared to fine-tuning-based methods (FT and GTOT), our proposed methods RCGrad and BLO+RCGrad demonstrate better performance across 7 out of 8 datasets. Specifically, compared to GTOT, both RCGrad and BLO+RCGrad achieve significant improvements of up to 5.1% and 1.8% in BBBP and ToxCast, respectively. Furthermore, RCGrad and BLO+RCGrad exhibit better performance than GS baselines with significantly improved ROC-AUC of as much as 7.7% and 2.2% in ClinTox and HIV, respectively. Overall, both RCGrad and BLO+RCGrad outperform fine-tuning methods, while achieving competitive or better performance than GS baselines across all datasets. Such consistently superior performance across multiple setups and pretrained GNNs underscores the robustness of our methods.

### Acknowledgements

The authors thank the anonymous reviewers for their constructive comments and feedback. The authors are grateful to the Ohio Supercomputer Center for their support on computational resources.

### Author contributions

XN conceived the research and supervised VD; VD and XN conducted the research; VD designed and implemented all methodologies; collected all materials, performed all the experiments and analyses, and prepared the initial draft of this manuscript; VD and XN reviewed and edited the manuscript.

### Funding

This project was made possible, in part, by support from the National Science Foundation grant no. IIS-2133650. Any opinions, findings, conclusions, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agency.

### Availability of data and materials

The data used in the manuscript are publicly and freely available from the MoleculeNet benchmark at <https://moleculenet.org/datasets-1>. All the code, materials, and instructions on data collection and code execution are publicly available at <https://github.com/vishaldehyiest/GraphTA>. All the required softwares and materials to execute the code are public and freely available.

## Declarations

### Competing interests

The authors declare that they have no competing interests.

Received: 18 April 2024 Accepted: 9 July 2024

Published online: 24 July 2024

## References

- (2022) Rdkit: Open-source cheminformatics. <http://www.rdkit.org>, accessed on May 2023
- Chen H, Wang X, Guan C et al (2022) Auxiliary learning with joint task and data scheduling. Proc 39th Int Machine Learn Proc Machine Learn Res 162:3634–3647
- Chen S, Crammer K, He H, et al (2021) Weighted training for cross-task learning. In: International Conference on Learning Representations
- Chen Z, Badrinarayanan V, Lee CY, et al (2018) GradNorm: Gradient normalization for adaptive loss balancing in deep multitask networks. International conference on machine learning, PMLR, pp 794–803
- David L, Thakkar A, Mercado R et al (2020) Molecular representations in ai-driven drug discovery: a review and practical guide. J Cheminform 12(1):1–22
- Dery LM, Dauphin Y, Grangier D (2021a) Auxiliary task update decomposition: The good, the bad and the neutral. arXiv preprint [arXiv:2108.11346](https://arxiv.org/abs/2108.11346)
- Dery LM, Michel P, Talwalkar A, et al (2021b) Should we be pre-training? an argument for end-task aware training as an alternative. In: International Conference on Learning Representations
- Dery LM, Michel P, Khodak M, et al (2022) Aang: Automating auxiliary learning. In: The Eleventh International Conference on Learning Representations
- Du Y, Czarnecki WM, Jayakumar SM, et al (2018) Adapting auxiliary losses using gradient similarity. arXiv preprint [arXiv:1812.02224](https://arxiv.org/abs/1812.02224)
- Floridi L, Chiriatti M (2020) Gpt-3: Its nature, scope, limits, and consequences. Minds Machines 30:681–694
- Gasteiger J, Yeshwanth C, Günnemann S (2021) Directional message passing on molecular graphs via synthetic coordinates. Adv Neural Inform Process Syst 34:15421–15433
- Guo Z, Guo K, Nan B et al (2023) Graph-based molecular representation learning In of the Thirty-Second International Joint Artificial Intelligence, IJCAI-23. Int Joint Conf Artif Intell Organ. <https://doi.org/10.24963/ijcai.2023/744>
- Gururangan S, Marasović A, Swayamdipta S, et al (2020) Don't stop pretraining: Adapt language models to domains and tasks. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp 8342–8360

14. Hamilton WL, Ying Z, Leskovec J (2017) Inductive Representation Learning on Large Graphs. In: NIPS, pp 1024–1034
15. He Y, Feng X, Cheng C et al (2022) Metabalance: improving multi-task recommendations via adapting gradient magnitudes of auxiliary tasks. *Proc ACM Web Conf 2022*:2205–2215
16. Houlsby N, Giurgiu A, Jastrzebski S et al (2019) Parameter-efficient transfer learning for NLP. *Proc 36th Int Conf machine Learn Res* 97:2790–2799
17. Hu W, Liu B, Gomes J, et al (2019) Strategies for pre-training graph neural networks. In: International Conference on Learning Representations
18. Hu Z, Dong Y, Wang K, et al (2020) Gpt-gnn: Generative pre-training of graph neural networks. In: Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp 1857–1867
19. Lee Y (2021) Improving end-to-end task-oriented dialog system with a simple auxiliary task. *Findings Assoc Computat Linguistics: EMNLP 2021*:1296–1303
20. Liebel L, Körner M (2018) Auxiliary tasks in multi-task learning. *arXiv preprint [arXiv:1805.06334](https://arxiv.org/abs/1805.06334)*
21. Liu S, Davison A, Johns E (2019a) Self-supervised generalisation with meta auxiliary learning. *Advances in Neural Information Processing Systems* 32
22. Liu S, Wang H, Liu W, et al (2021) Pre-training molecular graph representation with 3d geometry. In: International Conference on Learning Representations
23. Liu Y, Ott M, Goyal N, et al (2019b) Roberta: A robustly optimized bert pretraining approach. *arXiv preprint [arXiv:1907.11692](https://arxiv.org/abs/1907.11692)*
24. Lorraine J, Vicol P, Duvenaud D (2020) Optimizing millions of hyperparameters by implicit differentiation. In: International conference on artificial intelligence and statistics, PMLR, pp 1540–1552
25. Navon A, Achituve I, Maron H, et al (2020) Auxiliary learning by implicit differentiation. In: International Conference on Learning Representations
26. Nedyanchath A, Paramasivam P, Yenigalla P (2020) Multi-head attention for speech emotion recognition with auxiliary learning of gender recognition. *ICASSP 2020–2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, pp 7179–7183
27. Radford A, Narasimhan K, Salimans T, et al (2018) Improving language understanding by generative pre-training
28. Rong Y, Bian Y, Xu T et al (2020) Self-supervised graph transformer on large-scale molecular data. *Adv Neural Inform Process Syst* 33:12559–12571
29. Rosenstein MT, Marx Z, Kaelbling LP, et al (2005) To transfer or not to transfer. In: NIPS 2005 workshop on transfer learning
30. Ruder S (2017) An overview of multi-task learning in deep neural networks. *arXiv preprint [arXiv:1706.05098](https://arxiv.org/abs/1706.05098)*
31. Shamsian A, Navon A, Glazer N, et al (2023) Auxiliary learning as an asymmetric bargaining game. *arXiv preprint [arXiv:2301.13501](https://arxiv.org/abs/2301.13501)*
32. Shi B, Hoffman J, Saenko K et al (2020) Auxiliary task reweighting for minimum-data learning. *Adv Neural Inform Process Syst* 33:7148–7160
33. Sun FY, Hoffmann J, Verma V, et al (2019) Infograph: Unsupervised and semi-supervised graph-level representation learning via mutual information maximization. *arXiv preprint [arXiv:1908.01000](https://arxiv.org/abs/1908.01000)*
34. Sun R, Dai H, Yu AW (2022) Does gnn pretraining help molecular representation? *Adv Neural Inform Process Syst* 35:12096–12109
35. Trinh T, Dai A, Luong T, et al (2018) Learning longer-term dependencies in rnns with auxiliary losses. In: International Conference on Machine Learning, PMLR, pp 4965–4974
36. Wang H, Liu S, Kaddour J, et al (2022a) Evaluating self-supervised learned molecular graphs. In: ICML 2022 2nd AI for Science Workshop
37. Wang Y, Wang J, Cao Z et al (2022) Molecular contrastive learning of representations via graph neural networks. *Nature Machine Intell* 4(3):279–287
38. Wei J, Tay Y, Bommasani R, et al (2022) Emergent abilities of large language models. *arXiv preprint [arXiv:2206.07682](https://arxiv.org/abs/2206.07682)*
39. Wieder O, Kohlbacher S, Kuenemann M et al (2020) A compact review of molecular property prediction with graph neural networks. *Drug Discov Today Technol* 37:1–12
40. Wu Z, Ramsundar B, Feinberg EN et al (2018) Moleculenet: a benchmark for molecular machine learning. *Chem Sci* 9(2):513–530
41. Xia J, Zheng J, Tan C, et al (2022a) Towards effective and generalizable fine-tuning for pre-trained molecular graph models. *bioRxiv* pp 2022–02
42. Xia J, Zhu Y, Du Y, et al (2022b) Pre-training graph neural networks for molecular representations: retrospect and prospect. In: ICML 2022 2nd AI for Science Workshop
43. Xuhong L, Grandvalet Y, Davoine F (2018) Explicit inductive bias for transfer learning with convolutional networks. In: International Conference on Machine Learning, PMLR, pp 2825–2834
44. Yang E, Nair S, Chandradevan R, et al (2022) C3: Continued pretraining with contrastive weak supervision for cross language ad-hoc retrieval. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp 2507–2512
45. Yu T, Kumar S, Gupta A et al (2020) Gradient surgery for multi-task learning. *Adv Neural Inform Process Syst* 33:5824–5836
46. Zhang J, Xiao X, Huang LK, et al (2022) Fine-tuning graph neural networks via graph topology induced optimal transport. In: Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22. International Joint Conferences on Artificial Intelligence Organization, pp 3730–3736, <https://doi.org/10.24963/ijcai.2022/518>

## Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.