

---

# Precedence-Constrained Winter Value for Effective Graph Data Valuation

---

Hongliang Chi<sup>1</sup> Wei Jin<sup>2</sup> Charu Aggarwal<sup>3</sup> Yao Ma<sup>1</sup>

<sup>1</sup>Department of Computer Science, Rensselaer Polytechnic Institute, Troy, NY 12180

<sup>2</sup>Department of Computer Science, Emory University, Atlanta, GA 30322

<sup>3</sup>IBM T. J. Watson Research Center, Yorktown Heights, NY 10598

{chih3,may13}@rpi.edu, wei.jin@emory.edu, charu@us.ibm.com

## Abstract

1 Data valuation is essential for quantifying data’s worth, aiding in assessing data  
2 quality and determining fair compensation. While existing data valuation methods  
3 have proven effective in evaluating the value of Euclidean data, they face limita-  
4 tions when applied to the increasingly popular graph-structured data. Particularly,  
5 graph data valuation introduces unique challenges, primarily stemming from the  
6 intricate dependencies among nodes and the growth in value estimation costs. To  
7 address the challenging problem of graph data valuation, we put forth an innovative  
8 solution, **Precedence-Constrained Winter** (PC-Winter) Value, to account for the  
9 complex graph structure. Furthermore, we develop a variety of strategies to address  
10 the computational challenges and enable efficient approximation of PC-Winter.  
11 Extensive experiments demonstrate the effectiveness of PC-Winter across diverse  
12 datasets and tasks.

## 13 1 Introduction

14 The abundance of training data has been a key driver of recent advancements in machine learning  
15 (ML) [51]. As models and the requisite training data continue to expand in scale, data valuation  
16 has gained significant attention due to its ability to quantify the usefulness of data for ML tasks and  
17 determine fair compensation [28, 34]. Notable techniques in this field include Data Shapley [13]  
18 and its successors [20, 39, 29], which have gained prominence in assessing data value. Despite  
19 the promise of these methods, they are primarily designed for Euclidean data, where samples are  
20 often assumed to be independent and identically distributed (i.i.d.). Given the prevalence of graph-  
21 structured data in the real world [10, 31, 22], there arises a compelling need to perform data valuation  
22 for graphs. However, due to the interconnected nature of samples (nodes) on graphs, existing data  
23 valuation frameworks are not directly applicable to addressing the graph data valuation problem.

24 In particular, designing data valuation methods for graph-structured data faces several fundamen-  
25 tal challenges: **Challenge I:** Graph machine learning algorithms such as Graph Neural Networks  
26 (GNNs) [19, 37, 41] often involve both labeled and unlabeled nodes in their model training process.  
27 Therefore, unlabeled nodes, despite their absence of explicit labels, also hold intrinsic value. Existing  
28 data valuation methods, which typically assess a data point’s value based on its features and the  
29 associated label, do not readily accommodate the valuation of unlabeled nodes within graphs. **Chal-**  
30 **lenge II:** Nodes in a graph contribute to model performance in an interdependent and complex way:  
31 (1) Unlabeled nodes, while not providing direct supervision, can contribute to model performance by  
32 potentially affecting multiple labeled nodes through message-passing. (2) Labeled nodes, on the other

33 hand, contribute by providing direct supervision signals for model training, and similarly to unlabeled  
 34 nodes, they also contribute by affecting other labeled nodes through message-passing. **Challenge III:**  
 35 Traditional data valuation methods are often computationally expensive due to repeated retraining  
 36 of models [13]. The challenge is magnified in the context of graph-structured data, where samples  
 37 contribute to model performance in multifaceted manners. Additionally, the inherent message-passing  
 38 mechanism in GNN models further amplifies the computational demands for model re-training.

39 In this work, we make *the first attempt* to explore the challenging graph data valuation problem, to  
 40 the best of our knowledge. In light of the aforementioned challenges, we propose the **Precedence-**  
 41 **Constrained Winter (PC-Winter) Value**, a pioneering approach designed to intricately unravel and  
 42 analyze the contributions of nodes within graph structures, thereby offering a detailed perspective on  
 43 the valuation of graph elements. Our key contributions are as follows:

- 44 • We formulate the graph data valuation problem as a unique cooperative game [38] with special  
 45 coalition structures. Specifically, we decompose each node in the graph into several “players”  
 46 within the game, each representing a distinct contribution to model performance. We then devise  
 47 the PC-Winter to address the game, enabling the accurate valuation of all players. The PC-Winter  
 48 values of these players can be conveniently combined to generate values for nodes and edges.
- 49 • To tackle the computational challenges of calculating PC-Winter values, we develop a set of  
 50 strategies including *hierarchical truncation* and *local propagation*. These strategies together enable  
 51 an efficient approximation of PC-Winter values.
- 52 • Extensive experiments on various datasets and tasks, along with detailed ablation studies and  
 53 parameter analyses, validate the effectiveness of PC-Winter and provide insights into its behavior.

## 54 2 Preliminary and Related Work

55 In this section, we delve into some fundamental concepts that are essential for developing our  
 56 methodology. More extensive literature exploration can be found in Appendix A.

### 57 2.1 Cooperative Game Theory

58 Cooperative game theory explores the dynamics where players, or decision-makers, can form alliances,  
 59 known as coalitions, to achieve collectively beneficial outcomes [2, 7]. The critical components of  
 60 such a game include a *player set*  $\mathcal{P}$  consisting of all players in the game and a *utility function*  $U(\cdot)$ ,  
 61 which quantifies the value or payoff that each coalition of players can attain. Shapley Value [32] is  
 62 developed to fairly and efficiently distribute payoffs (values) among players.

63 **Shapley value.** The Shapley value  $\phi_i(\mathcal{P}, U)$  for a player  $i \in \mathcal{P}$  can be defined on permutations of  $\mathcal{P}$   
 64 as follows.


$$\phi_i(\mathcal{P}, U) = \frac{1}{|\Pi(\mathcal{P})|} \sum_{\pi \in \Pi(\mathcal{P})} [U(\mathcal{P}_i^\pi \cup \{i\}) - U(\mathcal{P}_i^\pi)] \quad (1)$$

65 where  $\Pi(\mathcal{P})$  denotes the set of all possible permutations of  $\mathcal{P}$  with  $|\Pi(\mathcal{P})|$  denoting its cardinality,  
 66 and  $\mathcal{P}_i^\pi$  is predecessor set of  $i$ , i.e, the set of players that appear before player  $i$  in a permutation  $\pi$ :

$$\mathcal{P}_i^\pi = \{j \in \mathcal{P} \mid \pi(j) < \pi(i)\}. \quad (2)$$

67 The Shapley value considers each player’s contribution to every possible coalition they could be  
 68 a part of. Specifically, in Eq. (1), for each permutation  $\pi$ , the *marginal contribution* of player  $i$  is  
 69 calculated as the difference in the utility function  $U$  when player  $i$  is added to an existing coalition  
 70  $\mathcal{P}_i^\pi$ . The Shapley value  $\phi_i(\mathcal{P}, U)$  for  $i$  is the average of these marginal contributions across all  
 71 permutations in  $\Pi(\mathcal{P})$ . The Shapley value has been widely applied in ML for various tasks such  
 72 as data valuation [13, 17] and model explanation [24, 11]. In the context of graph ML, it has been  
 73 primarily used for GNN explainability [8, 47, 25, 1]. A more detailed discussion on Shapley Value  
 74 on graph ML can be found in Appendix A.3.

75 **Winter Value.** The Shapley value is to address cooperative games, where  
 76 players collaborate freely and contribute on an equal footing. However, in many  
 77 practical cases, cooperative games, exhibit a *Level Coalition Structure* [26, 36,



78 48], reflecting a hierarchical organization. For instance, consider a corporate  
 79 setting where different tiers of management and staff contribute to a project in Coalition Structure

80 varying capacities and with differing degrees of decision-making authority. Players within such a game  
 81 are hierarchically categorized into nested coalitions with several levels, as depicted in Figure 1. The  
 82 outermost and largest ellipse represents the entire coalition and each of the smaller ellipse within the  
 83 largest ellipse symbolizes a “sub-coalition” at various hierarchical levels. Collaborations originate  
 84 within the smallest sub-coalitions at the base level (illustrated by the innermost ellipses in Figure 1.  
 85 These base units are then integrated into the next level, facilitating inter-coalition collaboration and  
 86 enabling contributions to ascend to higher levels. This bottom-up flow of contributions continues,  
 87 with each layer consolidating and passing on inputs to the next, culminating in a multi-leveled  
 88 collaborative contribution to the final objective of the entire coalition. To accommodate such  
 89 complex Level Coalition Structure, Winter value [40] was introduced. Winter value follows a similar  
 90 permutation-based definition as Shapley Value (Eq. (1)) but with only a specific subset of permutations  
 91 that respect the Level Coalition Structure. In these permutations, members of the same sub-coalition,  
 92 regardless of the level, must appear in an unbroken sequence without interruptions. This ensures that  
 93 the value attributed to each player is consistent with the level structure of the coalition. A formal  
 94 definition of the Winter value can be found in Appendix B.

## 95 2.2 Data Valuation and Data Shapley

96 Data valuation quantifies the contribution of data points for machine learning tasks. The seminal  
 97 work [13] introduces Data Shapley, applying cooperative game theory to data valuation, where  
 98 training samples are the *players*  $\mathcal{P}$ , and the *utility function*  $U$  assesses a model’s performance  
 99 on subsets of these players using a validation set. With  $\mathcal{P}$  and  $U$ , data values can be calculated  
 100 with Eq. (1). However, Data Shapley and subsequent methods [13, 20, 39] primarily focus on i.i.d.  
 101 data, overlooking potential coalitions or dependencies among data points.

## 102 2.3 Graphs and Graph Neural Networks

103 Consider a graph  $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$  where  $\mathcal{V}$  denotes the set of nodes and  $\mathcal{E}$  denotes the set of edges. Each  
 104 node  $v_i \in \mathcal{V}$  carries a feature vector  $\mathbf{x}_i \in \mathbb{R}^d$ , where  $d$  is the dimensionality of the feature space.  
 105 Additionally, each node  $v_i$  is associated with a label  $y_i$  from a set of possible labels  $\mathcal{C}$ . We assume  
 106 that only a subset  $\mathcal{V}_l \subset \mathcal{V}$  are with known labels.

107 GNNs [19, 37, 41] are prominent models for graph ML tasks. Specifically, from a local per-  
 108 spective for node  $v_i$ , the  $k$ -th GNN layer generally performs a feature averaging process as  
 109  $\mathbf{h}_i^{(k)} = \frac{1}{deg(v_i)} \sum_{v_j \in \mathcal{N}(v_i)} \mathbf{W} \mathbf{h}_j^{(k-1)}$ , where  $\mathbf{W}$  is the parameter matrix,  $deg(v_i)$  and  $\mathcal{N}(v_i)$  denote

110 the degree and neighbors of node  $v_i$ , respectively. After a total of  $K$  layers,  $\mathbf{h}_i^{(K)}$  are utilized as the  
 111 learned representation of  $v_i$ . Such a feature aggregation process can be also described with a  $K$ -level  
 112 *computation tree* [15] rooted on node  $v_i$ .

113 **Definition 1** (Computation Tree). *For a node  $v_i \in \mathcal{V}$ , its  $K$ -level computation tree corresponding to*  
 114 *a  $K$ -layer GNN model is denoted as  $\mathcal{T}_i^K$  with  $v_i$  as its root node. The first level of the tree consists of*  
 115 *the immediate neighbors of  $v_i$ , and each subsequent level is formed by the neighbors of nodes in the*  
 116 *level directly above. This pattern of branching out continues, expanding through successive levels of*  
 117 *neighboring nodes until the depth of the tree grows to  $K$ .*

118 The feature aggregation process in a  $K$ -layer GNN can be regarded as a bottom-up feature propagation  
 119 process in the computation tree, where nodes in the lowest level are associated with their initial  
 120 features. Therefore, the final representation  $\mathbf{h}_i^{(K)}$  of a node  $v_i$  is affected by all nodes within its  
 121  $K$ -hop neighborhood, which is referred to as the *receptive field* of node  $v_i$ . The GNN model is trained  
 122 using the  $(\mathbf{h}_i^{(K)}, y_i)$  pairs, where each labeled node  $v_i$  in  $\mathcal{V}_l$  is represented by its final representation  
 123 and corresponding label. *Thus, in addition to labeled nodes, those unlabeled nodes that are within*  
 124 *the receptive field of labeled nodes also contribute to model performance.*

## 125 3 Methodology

126 In classic machine learning models designed for Euclidean data, such as images and texts, training  
 127 samples are typically assumed as i.i.d. Thus, each labeled sample contributes to the model perfor-  
 128 mance by directly providing supervision signals through the training objective. However, due to the

interdependent nature of graph data, nodes in a graph contribute to GNN performance in a more complicated way, which poses unique challenges. Specifically, as discussed in Section 2.3, both labeled and unlabeled nodes are involved in the training stage through the feature aggregation. Next, we discuss how these nodes contribute to GNN performance.

**Observation 1.** *Unlabeled nodes influence GNN performance by affecting the final representation of labeled nodes. On the other hand, labeled nodes can contribute to GNN performance in two ways: (1) they provide direct supervision signals to GNN with their labels, and (2) just like unlabeled nodes, they can impact the final representation of other labeled nodes through feature aggregation. Note that both labeled nodes and unlabeled nodes can affect the final representations of multiple labeled nodes, as long as they lie within the receptive field of these labeled nodes. Hence, a single node can make multifaceted and heterogeneous contributions to GNN performance by affecting multiple labeled nodes in various manners.*

### 3.1 The Graph Data Valuation Problem

Based on Observation 1, due to the heterogeneous and diverse effects of labeled and unlabeled nodes, it is necessary to perform fine-grained data valuation on graph data elements. In particular, we propose to decompose a node into distinct “duplicates” corresponding to their impact on different labeled nodes. We then aim to obtain values for all “duplicates” of these nodes. This could clearly express and separate how nodes impact GNN performance in various aspects. Following existing literature [13, 39, 43], we approach the graph data value problem through a cooperative game. Next, we introduce the *player set* and the *utility function* of this game. In general, we define the graph data valuation game based on  $K$ -layer GNN models.

**Definition 2 (Player Set).** *The player set  $\mathcal{P}$  in a graph data valuation game is defined as the union of nodes in the computation trees of labeled nodes. Duplication of nodes may occur within a single computation tree  $\mathcal{T}_i^K$  or across different labeled nodes’ computation trees. In the graph data valuation game, these potential duplicates are treated as distinct players, uniquely identified by their paths to the corresponding labeled node. We define the player set  $\mathcal{P}$  as the set of all these distinct players across the computation trees of all labeled nodes in  $\mathcal{V}_l$ .*

**Definition 3 (Utility Function).** *Given a subset  $\mathcal{S} \subset \mathcal{P}$ , we first generate a node-induced graph  $G_{in}(\mathcal{S})$  using their corresponding edges in the computation trees. Then, a GNN model  $\mathcal{A}$  is trained on the induced graph  $G_{in}(\mathcal{S})$ . Its performance is evaluated on a held-out validation set to serve as the utility of  $\mathcal{S}$ , calculated as  $U(\mathcal{S}) = acc(\mathcal{A}(G_{in}(\mathcal{S})))$ , where  $acc$  measures the accuracy of the trained GNN model  $\mathcal{A}(G_{in}(\mathcal{S}))$  on a held-out validation set.*

The goal of the graph data valuation problem is to assign a value to all players in  $\mathcal{P}$  with the help of the *utility function*  $U$ . When calculated properly, these values are supposed to provide a detailed understanding of how players in  $\mathcal{P}$  contribute to the GNN performance in a fine-grained manner. Furthermore, these values can be flexibly combined to generate higher-level values for nodes and edges, which will be discussed in Section 3.5.


### 3.2 Precedence-Constrained Winter Value

As discussed in Section 2.3, the final representations of a labeled node  $v_i$  come from the hierarchical collaboration of all players in the computation tree  $\mathcal{T}_i^K$ . These labeled nodes with the updated representations then contribute to the GNN performance through the training objective. Such a contribution process forms a hierarchical collaboration between the players in  $\mathcal{P}$ , which can be illustrated with a *contribution tree*  $\mathcal{T}$  as shown in Figure 2a. In particular, the *contribution tree*  $\mathcal{T}$  is constructed by linking the root nodes of the computation trees of all labeled nodes with a dummy node representing the GNN training objective  $\mathcal{O}$ . In Figure 2a, for the ease of illustration, we set  $K = 2$ , include only 2 labeled nodes, i.e.  $v_0, v_1$ , and utilize  $w_i, u_i$  to denote the nodes in the lower level. The subtree rooted at a labeled node  $v_i \in \mathcal{V}$  is the corresponding computation tree  $\mathcal{T}_i^2$ . With this, we observe the following about the coalition structure of the graph data valuation game.

**Observation 2 (Level Coalition Structure).** *As shown in Figure 2a, the players in  $\mathcal{P}$  hierarchically collaborate to contribute. At the bottom level, the players are naturally grouped by their parents. Specifically, players with a common parent such as  $u_0, u_1, u_2$  with their parent  $w_0$ , establish a foundation sub-coalition. This sub-coalition is clearly depicted in Figure 2b. Moving up the tree,*

181 these parent nodes, like  $w_0$ , serve as “delegates” for their respective sub-coalitions, further engaging  
 182 in collaborations with other sub-coalitions. This interaction forms higher-level sub-coalitions, such  
 183 as the one between  $w_0$ ,  $w_1$ ,  $w_2$ , and  $v_0$  in Figure 2b, indicating inter-coalition cooperation. This  
 184 ascending process of coalition formation continues until the root node  $\mathcal{O}$  is reached, which represents  
 185 the objective of the entire coalition consisting of all players. The depicted hierarchical collaboration  
 186 process aligns with the Level Coalition Structure discussed in Section 2.1.

187 While the contribution tree shares similarities with the  
 188 Level Coalition Structure illustrated in Section 2.1, a  
 189 pivotal distinction lies in the representation and func-  
 190 tion of “delegates” (highlighted in red in Figure 2b)  
 191 within each coalition. In the traditional Level Coali-  
 192 tion Structure, contributions within a sub-coalition are  
 193 made collectively, with each player or lower-level sub-  
 194 coalition participating on an equitable basis. In contrast, the contribution tree framework distinguishes  
 195 itself by designating a “delegate” within each sub-coalition, a player that represents and advances  
 196 the collective contributions, establishing a directed and tiered flow of influence, hence forming a  
 197 Unilateral Dependency Structure.



198 **Observation 3** (Unilateral Dependency Structure). *In the contribution tree framework, a player*  
 199  *$p \in \mathcal{P}$  contributes to the final objective through a hierarchical pathway facilitated by its ancestors*  
 200 *(its “delegates” at different levels). Therefore, the collaboration between players in  $\mathcal{P}$  exhibits a*  
 201 *Unilateral Dependency Structure, where a player  $p$ ’s contribution is dependent on its ancestors.*

202 According to these two observations, the players demonstrate unique coalition structures in the graph  
 203 data valuation game. We aim to propose a permutation-based valuation framework similar to Eq. (1)  
 204 to address the cooperative game with both Level Coalition Structure and Unilateral Dependency  
 205 Structure. In particular, instead of utilizing all the permutations as in Eq. (1), only the *permissible*  
 206 *permutations* aligning with such coalition structures are included in the value calculations. As we  
 207 described in Section 2.1, cooperative games with Level Coalition Structure have been addressed by  
 208 the Winter value [40, 4]. Specifically, a permutation respecting the Level Coalition Structure must  
 209 ensure that players in the same (sub-)coalition, regardless of its level, are grouped together without  
 210 interruption from other players [40]. In our scenario, any subtree of the contribution tree corresponds  
 211 to a sub-coalition as demonstrated in Figure 2. Hence, we need to ensure that for any player  $p \in \mathcal{P}$ ,  
 212 the player  $p$  and its descendants in the contribution tree should be grouped together in the permutation.  
 213 For example, the players  $w_0$ ,  $u_0$ ,  $u_1$ ,  $u_2$  should present together as a group in the permutation with  
 214 potentially different orders. On the other hand, to ensure the Unilateral Dependency Structure, a  
 215 permutation must maintain a partial order. Specifically, for any player  $p$  in the permutation, its  
 216 descendants must present in later positions in the permutation than  $p$ . Otherwise, the descendants of  
 217  $p$  cannot make non-trivial contributions, resulting in 0 marginal contributions.

218 We formally define the *permissible permutations* that align with both Level Coalition Structure and  
 219 Unilateral Dependency Structure utilizing the following two constraints.

220 **Constraint 1** (Level Constraint). *For any player  $p \in \mathcal{P}$ , the set of its descendants in the contribution*  
 221 *tree is denoted as  $\mathcal{D}(p)$ . Then, a permutation  $\pi$  aligning with the Level Coalition Structure satisfies*  
 222 *the following Level Constraint:  $|\pi[i] - \pi[j]| \leq |\mathcal{D}(p)|, \forall i, j \in \mathcal{D}(p) \cup p, \forall p \in \mathcal{P}$ , where  $\pi[i]$  denotes*  
 223 *the positional rank of the  $i$  in  $\pi$ .*

224 **Constraint 2** (Precedence Constraint). *A permutation  $\pi$  aligning with the Unilateral Dependency*  
 225 *Structure satisfies the following Precedence Constraint:  $\pi[p] < \pi[i], \forall i \in \mathcal{D}(p), \forall p \in \mathcal{P}$ .*

226 We denote the set of *permissible permutations* satisfying both *Level Constraint* and *Precedence*  
 227 *Constraint* as  $\Omega$ . Then, we define the **Precedence-Constrained Winter** (PC-Winter) value for a  
 228 player  $p \in \mathcal{P}$  with the permutations in  $\Omega$  as follows.

$$\psi_p(\mathcal{P}, U) = \frac{1}{|\Omega|} \sum_{\pi \in \Omega} (U(\mathcal{P}_p^\pi \cup p) - U(\mathcal{P}_p^\pi)), \quad (3)$$

229 where  $U(\cdot)$  is the utility function (see Definition 3), and  $\mathcal{P}_p^\pi$  denotes the predecessor set of  $p$  in  $\pi$  as  
 230 defined in Eq. (2).

### 231 3.3 Permissible Permutations for PC-Winter

232 To calculate PC-Winter value, it is required to obtain all permissible permutations. A straightforward  
233 way is to enumerate all permutations and only retain the permissible permutations. However, such an  
234 approach is extremely computationally intensive and typically not feasible in reality. In this section,  
235 to address this challenge, we propose a novel method to directly generate these permutations by  
236 traversing the contribution tree with Depth-First Search (DFS). Specifically, each DFS traversal results  
237 in a *preordering*, which is a list of the nodes (players) in the order that they were visited by DFS.  
238 Such a *preordering* naturally defines a permutation of  $\mathcal{P}$  by simply removing the dummy node in the  
239 contribution tree from the *preordering*. By iterating all possible DFS traversals of the contribution  
240 tree, we can obtain all permutations in  $\Omega$ , which is demonstrated in the following theorems.

241 **Theorem 1** (Specificity). *Given a contribution tree  $\mathcal{T}$  with a set of players  $\mathcal{P}$ , any DFS traversal*  
242 *over the  $\mathcal{T}$  results in a permissible permutation of  $\mathcal{P}$  that satisfies both the Level Constraint and*  
243 *Precedence Constraint.*

244 **Theorem 2** (Exhaustiveness). *Given a contribution tree  $\mathcal{T}$  with a set of players  $\mathcal{P}$ , any permissible*  
245 *permutation  $\pi \in \Omega$  can be generated by a corresponding DFS traversal of  $\mathcal{T}$ .*

246 The proofs for two theorems can be found in Appendix C. Theorem 1 demonstrates that DFS  
247 traversals *specifically* generate *permissible permutations*. On the other hand, Theorem 2 ensures  
248 the *exhaustiveness* of generation, which allows us to obtain all permutations in  $\Omega$  by DFS traversal.  
249 Together, these two theorems ensure us to *exactly* generate the set of *permissible permutations*  $\Omega$ .

250 Notably, the calculation of PC-Winter value involves two steps: 1) generating  $\Omega$  with DFS traversals;  
251 and 2) calculating the PC-Winter value according to Eq. (3). Nonetheless, it can be done in a  
252 streaming way while we perform the DFS traversals. Specifically, once we reach a player  $p$  in a DFS  
253 traversal, we can immediately calculate its marginal contribution. The PC-Winter values for all  
254 players are computed by averaging their marginal contributions from all possible DFS traversals.

### 255 3.4 Efficient Approximation of PC-Winter

256 Calculating the PC-Winter value for players in  $\mathcal{P}$  is infeasible due to computational intensity, arising  
257 from: 1) The exponential growth in the number of permissible permutations with more players,  
258 rendering exhaustive enumeration intractable; 2) The necessity to re-train the GNN within the utility  
259 function for each permutation, a process repeated  $|\mathcal{P}|$  times to account for every player’s marginal  
260 contribution; and 3) The intensive computation involved in GNN re-training, requiring feature  
261 aggregation over the graph that increases in complexity with the graph’s size. These challenges  
262 necessitate an efficient approximation method for PC-Winter valuation in practical applications. We  
263 propose three strategies to address these computational issues.

264 **Permutation Sampling.** Following Data Shapley [13], we adopt Monte Carlo (MC) sampling to  
265 randomly sample a subset of permissible permutations denoted as  $\Omega_s$ . Then, we utilize  $\Omega_s$  to replace  
266  $\Omega$  in Eq. (3) for approximating PC-Winter value.

267 **Hierarchical Truncation.** GNN models often demonstrate a phenomenon of *neighborhood saturation*,  
268 i.e., these models achieve satisfactory performance even when trained on a subgraph using only  
269 a small subset of neighbors, rather than the full neighborhood [14, 23, 45, 5], indicating diminishing  
270 returns from additional neighbors beyond a certain point. This indicates that for a player  $p$  in a  
271 permissible permutation  $\pi$  generated by DFS over the contribution tree, the marginal contributions of  
272 its late visited child players are insignificant. Thus, we propose hierarchical truncation for efficiently  
273 obtaining the marginal contributions by directly approximating insignificant values as 0. Specifically,  
274 during the DFS traversal, given a truncation ratio  $r$ , we only compute actual marginal contributions  
275 for players in the first  $1 - r$  portion of each node’s child subtrees, approximating the marginal  
276 contributions of players in the remaining subtrees as 0. For example, in Figure 2a, given a truncation  
277 ratio  $r = 2/3$ , when DFS reaches player  $v_0$ , we only calculate marginal contributions for players in  
278 the subtree rooted at  $w_0$ . Furthermore, in the subtree rooted at  $w_0$ , due to the hierarchical truncation,  
279 only the marginal contribution of  $u_0$  is evaluated, those for node  $u_1$  and  $u_2$  are set to 0. This approach  
280 is further optimized by adjusting truncation ratios based on the tree level, accommodating varying  
281 contribution patterns across levels. In particular, we organize the pair of truncation ratio as  $r_1-r_2$ ,

282 indicating we truncate  $r_1$  (or  $r_2$ ) portion of subtrees (or child players) of  $v_i$  (or  $w_i$ ). We show how  
283 the hierarchical truncation helps tremendously reduce the model re-training in Appendix D.

284 **Local Propagation.** To enhance scalability, we leverage SGC [41] in our utility function, which  
285 simplifies GNNs by aggregating node features before applying an MLP. According to the Level  
286 Constraint (Constraint 1), the players within the same computation tree are grouped together in the  
287 permutation. Therefore, the induced graph of any coalition  $\mathcal{P}_p^\pi$  defined by a permissible permutation  
288 consists of a set of separated computation trees (or a partial computation tree corresponding to the last  
289 visited labeled node in  $\mathcal{P}_p^\pi$ ). A key observation is that the feature aggregation process for the labeled  
290 nodes can be done independently within their own computation trees. Hence, instead of performing  
291 the feature propagation for the entire induced graph, we propose to perform *local propagation* only  
292 on necessary computation trees. In particular, the aggregated representation for a labeled node is  
293 fixed after we traverse its entire computation tree in DFS. Therefore, for evaluating a player  $p$ 's  
294 marginal contribution, only the partial computation tree of the last visited labeled node requires *local*  
295 *propagation*, minimizing feature propagation efforts.

296 The PC-Winter values for all players are approximated with these three strategies in a streaming  
297 manner. In particular, we randomly traverse the contribution tree with DFS for  $|\Omega_s|$  times. During  
298 each DFS traversal, the marginal contributions for all players in  $\mathcal{P}$  are efficiently obtained with the  
299 help of *hierarchical truncation* and *local propagation*. The marginal contributions calculated through  
300 these  $|\Omega_s|$  DFS traversals are averaged to approximate the PC-Winter value for all players. In  
301 Appendix H.5, we provide a detailed complexity analysis of the PC-Winter algorithm.

### 302 3.5 From PC-Winter to Node and Edge Values

303 The PC-Winter values for players in  $\mathcal{P}$  can be flexibly combined to obtain the values for elements  
304 in the original graph, which are illustrated in this section. Specifically, as discussed in Section 3.1,  
305 multiple “duplicates” of a node  $v \in \mathcal{V}$  in the original graph may potentially present in  $\mathcal{P}$ . Thus, we  
306 could obtain *node value* for the node  $v$  by summing the PC-Winter values of all its “duplicates” in  $\mathcal{P}$ .  
307 On the other hand, each player (except for the rooted labeled players) in  $\mathcal{P}$  corresponds to an “edge” in  
308 the contribution tree as identified by the player and its parent. For instance, in Figure 2a, the player  $u_0$   
309 corresponds to “edge” connecting  $u_0$  and  $w_0$ . Therefore, DFS traversals also generate permutations  
310 for these “edges”. From this perspective, the marginal contribution for a player  $p$  calculated through  
311 a DFS traversal can be also regarded as the marginal contribution of its corresponding edge, if we  
312 treat this process as gradually adding “edges” to connecting the players in  $\mathcal{P}$ . Hence, the PC-Winter  
313 values for players in  $\mathcal{P}$  can be regarded as PC-Winter values for their corresponding “edges” in the  
314 contribution tree. Multiple “duplicates” of an edge  $e \in \mathcal{E}$  in the original graph may be present in the  
315 contribution tree. Hence, similar to the *node values*, we define the *edge value* for  $e \in \mathcal{E}$  by taking the  
316 summation of the PC-Winter value for all its “duplicates” in the contribution tree.

## 317 4 Experiment

318 **Datasets and Settings.** We assess the proposed approach on six real-world benchmark datasets:  
319 Cora, Citeseer, and Pubmed [30], Amazon-Photo, Amazon-Computer, and Coauthor-Physics [33].  
320 The detailed statistics of datasets are summarized in Table 2 in Appendix G. Our experiments focus  
321 on the inductive node classification task. The detailed setup of the inductive setting can be found in  
322 Appendix G.1. To obtain the PC-Winter values, we run permutations in a streaming way as described  
323 in Section 3.4. This process terminates with a convergence criterion as detailed in Appendix G.4.  
324 PC-Winter typically terminates with a different number of permutations for different datasets. The  
325 other hyper-parameters are detailed in Appendix G.5.

### 326 4.1 Dropping High-Value Nodes

327 In this section, we aim to evaluate the quality of data values produced by PC-Winter via dropping  
328 high-value nodes from the graph. Dropping high-value nodes is expected to significantly diminish  
329 performance, and thus the performance observed after removing high-value nodes serves as a strong  
330 indicator of the efficacy of graph data valuation. Notably, PC-Winter values are calculated as  
331 described in Section 3.5.

332 To demonstrate the effectiveness of PC-Winter, we include Random value, Degree value, Leave-  
 333 one-out (LOO) value, and Data Shapley value as baselines. A more detailed description of these  
 334 baselines is included in Appendix G.6. To conduct node-dropping experiments, nodes are ranked by  
 335 their assessed values for each method and removed sequentially from the training graph  $\mathcal{G}_{tr}$ . After  
 336 each removal, we train a GNN model based on the remaining graph and evaluate its performance  
 337 on the testing graph  $\mathcal{G}_{te}$ . Performance changes are depicted through a curve that tracks the model’s  
 338 accuracy as nodes are progressively eliminated. Labeled nodes often contribute more significantly  
 339 to model performance than unlabeled nodes because they directly offer supervision. Thereby, with  
 340 accurately assigned node values, labeled nodes should be prioritized for removal over unlabeled  
 341 nodes. We empirically validate this hypothesis in Figure 6, discussed in Appendix E. Specifically, in  
 342 nearly all datasets, our observations reveal that the majority of labeled nodes are removed prior to the  
 343 unlabeled nodes by both PC-Winter and Data Shapley. This leads to a plateau in the latter portion  
 344 of the performance curves since a GNN model cannot be effectively trained with only unlabeled  
 345 nodes. Consequently, this scenario significantly hampers the ability to assess the value of unlabeled  
 346 nodes. Therefore, we propose to conduct separate assessments for the values of labeled and unlabeled  
 347 nodes. Here, we only include the results for unlabeled nodes, while the results for labeled nodes are  
 348 presented in Appendix F.

349 **Results and Analysis.** Figure 3 illustrates the performance comparison between PC-Winter and  
 350 other baselines across various datasets. From Figure 3, we make the following observations. First,  
 351 the removal of high-value unlabeled nodes identified by PC-Winter consistently results in the  
 352 most significant decline in model performance across various datasets. This is particularly evident  
 353 after removing a relatively small fraction (10%-20%) of the highest-value nodes. This trend  
 354 underscores the importance of high-value nodes. Notably, in most datasets PC-Winter outper-  
 355 forms the best baseline method, Data Shapley, by a considerable margin, highlighting its effective-  
 356 ness. Second, the decrease in performance caused by our method is not only substantial  
 357 but also persistent throughout the node-dropping process, further validating the effectiveness of  
 358 PC-Winter. Third, the performance curves of PC-Winter and Data Shapley eventually rebound  
 359 towards the end. This rebound corresponds to the removal of unlabeled nodes that make negative con-  
 360 tributions. Their removal aids in improving performance, ultimately reaching the MLP performance  
 361 when all nodes are excluded. This upswing not only evidences the discernment of PC-Winter and  
 362 Data Shapley in ascertaining node values but also showcases the particularly acute precision of  
 363 PC-Winter. These insights collectively affirm the capability of PC-Winter in accurately assessing  
 364 node values.




Figure 3: Dropping High-Value Nodes

#### 374 4.2 Adding High-Value Edges

375 In this section, we explore the impact of adding high-value elements to a graph, providing an  
 376 alternative perspective to validate the effectiveness of data valuation. Notably, adding high-value  
 377 nodes to a graph typically involves the concurrent addition of edges, which complicates the addition  
 378 process. Thus, we target the addition of high-value edges, providing a complementary perspective  
 379 to our analysis. As described in Section 3.5, the flexibility of PC-Winter allows for obtaining edge  
 380 values without a separate “reevaluation” process for edges.

381 Here, we keep all nodes in  $\mathcal{G}_{tr}$  and sequentially add edges according to the edge values in descend-  
 382 ing order, starting with the highest-valued ones. Similar to the node-dropping experiments, the  
 383 effectiveness of the edge addition is shown through performance curves. We include Random value,



384 Edge-Betweenness, Leave-one-out (L00) as baselines. Notably, here, Random and L00 specifically  
 385 pertain to edges, and while we use the same terminology as in the prior section, they are distinct  
 386 methods, which are detailed in Appendix G.6.

387 **Results and Analysis.** Figure 4 illustrates that  
 388 the Random, L00, and Edge-Betweenness base-  
 389 lines achieve only linear performance improve-  
 390 ments with the addition of more edges, failing  
 391 to discern the most impactful ones for a sparse  
 392 yet informative graph. In contrast, the inclu-  
 393 sion of edges based on the PC-Winter value re-  
 394 sults in a steep performance climb, affirming the  
 395 PC-Winter’s efficacy in pinpointing key edges.  
 396 Notably, the Cora dataset reaches full-graph per-  
 397 formance using merely 8% of the edges selected  
 398 by PC-Winter. Moreover, with just 10% of  
 399 PC-Winter-selected edges, the accuracy climbs  
 400 to 72.9%, outperforming the full graph’s 71.3%,  
 401 underscoring PC-Winter’s capability to identify  
 402 valuable edges. This trend is generally consistent  
 403 across other datasets as well.

### 4.3 Ablation Study, Parameter and Efficiency Analysis

404 In this section, we conduct an ablation study, parameter analysis, and efficiency analysis to gain  
 405 deeper insights into PC-Winter using node-dropping experiments.  
 406

407 **Ablation Study.** We conduct an ablation study to  
 408 understand how the two constraints in Section 3.2  
 409 affect the effectiveness of PC-Winter. We intro-  
 410 duce two variants of PC-Winter by lifting one  
 411 of the constraints for the permutations. In par-  
 412 ticular, we define PC-Winter-L using the per-  
 413 mutations satisfying the Level Constraint. Simi-  
 414 larly, PC-Winter-P is defined with permutations  
 415 only satisfying Precedence Constraint. As shown in Figure 5, PC-Winter value outperforms the  
 416 PC-Winter-L and PC-Winter-P on both datasets, which demonstrates that both constraints are  
 417 crucial for PC-Winter. Additional results on other datasets are provided in Appendix H.1.

418 **Parameter Analysis.** We conduct parameter analyses to investigate the impact of permutation  
 419 number and truncation ratios on PC-Winter’s performance. The results reveal that PC-Winter  
 420 achieves robust performance even with a significantly reduced number of permutations and high  
 421 truncation ratios. Detailed findings are presented in Appendix H.2 and Appendix H.3, respectively.

422 **Efficiency Analysis.** We compare the efficiency of PC-Winter and Data Shapley. Analysis of  
 423 converged permutation count and time per permutation across 6 datasets underscores PC-Winter’s  
 424 significantly higher efficiency. A comprehensive breakdown is available in Appendix H.4.

## 5 Conclusion

425  
 426 In this paper, we introduce PC-Winter, an innovative approach for effective graph data valuation.  
 427 The method is specifically designed for graph-structured data and addresses the challenges posed by  
 428 unlabeled elements and complex node dependencies within graphs. Furthermore, we introduce a set  
 429 of strategies for reducing the computational cost, enabling efficient approximation of PC-Winter.  
 430 Extensive experiments demonstrate the practicality and effectiveness of PC-Winter in various  
 431 datasets and tasks. While PC-Winter demonstrates improved efficiency compared to Data Shapley,  
 432 we acknowledge that further efficiency enhancements are crucial to fully unlock the potential of graph  
 433 data valuation in real-world applications. Our work can be seen as a foundation for future research in  
 434 this direction.




Figure 4: Adding the High-Value Edges




Figure 5: Ablation Study

## 435 References

- 436 [1] Selahattin Akkas and Ariful Azad. Gnnshap: Scalable and accurate gnn explanation using  
437 shapley values. In *Proceedings of the ACM on Web Conference 2024*, pages 827–838, 2024.
- 438 [2] Rodica Branzei, Dinko Dimitrov, and Stef Tijs. *Models in cooperative game theory*, volume  
439 556. Springer Science & Business Media, 2008.
- 440 [3] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally  
441 connected networks on graphs. *arXiv preprint arXiv:1312.6203*, 2013.
- 442 [4] Frédéric Chantreuil. Axiomatics of level structure values. *Power Indices and Coalition*  
443 *Formation*, pages 45–62, 2001.
- 444 [5] Jianfei Chen, Jun Zhu, and Le Song. Stochastic training of graph convolutional networks with  
445 variance reduction. *arXiv preprint arXiv:1710.10568*, 2017.
- 446 [6] Zizhang Chen, Peizhao Li, Hongfu Liu, and Pengyu Hong. Characterizing the influence of  
447 graph elements. *arXiv preprint arXiv:2210.07441*, 2022.
- 448 [7] Imma Curiel. *Cooperative game theory and applications: cooperative games arising from*  
449 *combinatorial optimization problems*, volume 16. Springer Science & Business Media, 2013.
- 450 [8] Alexandre Duval and Fragkiskos D Malliaros. Graphsvx: Shapley value explanations for graph  
451 neural networks. In *Machine Learning and Knowledge Discovery in Databases. Research Track:*  
452 *European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021, Proceedings,*  
453 *Part II 21*, pages 302–318. Springer, 2021.
- 454 [9] Edoardo D’Amico, Khalil Muhammad, Elias Tragos, Barry Smyth, Neil Hurley, and Aonghus  
455 Lawlor. Item graph convolution collaborative filtering for inductive recommendations. In  
456 *European Conference on Information Retrieval*, pages 249–263. Springer, 2023.
- 457 [10] Wenqi Fan, Yao Ma, Dawei Yin, Jianping Wang, Jiliang Tang, and Qing Li. Deep social  
458 collaborative filtering. In *Proceedings of the 13th ACM RecSys*, 2019.
- 459 [11] Christopher Frye, Damien de Mijolla, Tom Begley, Laurence Cowton, Megan Stanley, and Ilya  
460 Feige. Shapley explainability on the data manifold. *arXiv preprint arXiv:2006.01272*, 2020.
- 461 [12] Johannes Gasteiger, Aleksandar Bojchevski, and Stephan Günnemann. Predict then propagate:  
462 Graph neural networks meet personalized pagerank. *arXiv preprint arXiv:1810.05997*, 2018.
- 463 [13] Amirata Ghorbani and James Zou. Data shapley: Equitable valuation of data for machine  
464 learning. In *International conference on machine learning*, pages 2242–2251. PMLR, 2019.
- 465 [14] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large  
466 graphs. *Advances in neural information processing systems*, 30, 2017.
- 467 [15] Stefanie Jegelka. Theory of graph neural networks: Representation and learning. *arXiv preprint*  
468 *arXiv:2204.07697*, 2022.
- 469 [16] Theis E Jendal, Matteo Lissandrini, Peter Dolog, and Katja Hose. Simple and powerful  
470 architecture for inductive recommendation using knowledge graph convolutions. *arXiv preprint*  
471 *arXiv:2209.04185*, 2022.
- 472 [17] Ruoxi Jia, David Dao, Boxin Wang, Frances Ann Hubis, Nezihe Merve Gurel, Bo Li, Ce Zhang,  
473 Costas J Spanos, and Dawn Song. Efficient task-specific data valuation for nearest neighbor  
474 algorithms. *arXiv preprint arXiv:1908.08619*, 2019.
- 475 [18] Hoang Anh Just, Feiyang Kang, Jiachen T Wang, Yi Zeng, Myeongseob Ko, Ming Jin, and  
476 Ruoxi Jia. Lava: Data valuation without pre-specified learning algorithms. *arXiv preprint*  
477 *arXiv:2305.00054*, 2023.

- 478 [19] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional  
479 networks. *arXiv preprint arXiv:1609.02907*, 2016.
- 480 [20] Yongchan Kwon and James Zou. Beta shapley: a unified and noise-reduced data valuation  
481 framework for machine learning. *arXiv preprint arXiv:2110.14049*, 2021.
- 482 [21] Yongchan Kwon and James Zou. Data-oob: Out-of-bag estimate as a simple and efficient data  
483 value. *arXiv preprint arXiv:2304.07718*, 2023.
- 484 [22] Pengyong Li, Jun Wang, Yixuan Qiao, Hao Chen, Yihuan Yu, Xiaojun Yao, Peng Gao, Guotong  
485 Xie, and Sen Song. An effective self-supervised framework for learning expressive molecular  
486 global representations to drug discovery. *Briefings in Bioinformatics*.
- 487 [23] Xin Liu, Mingyu Yan, Lei Deng, Guoqi Li, Xiaochun Ye, and Dongrui Fan. Sampling meth-  
488 ods for efficient training of graph convolutional networks: A survey. *IEEE/CAA Journal of*  
489 *Automatica Sinica*, 9(2):205–234, 2021.
- 490 [24] Zelei Liu, Yuanyuan Chen, Han Yu, Yang Liu, and Lizhen Cui. Gtg-shapley: Efficient and accu-  
491 rate participant contribution evaluation in federated learning. *ACM Transactions on Intelligent*  
492 *Systems and Technology (TIST)*, 13(4):1–21, 2022.
- 493 [25] Andrea Mastropietro, Giuseppe Pasculli, Christian Feldmann, Raquel Rodríguez-Pérez, and  
494 Jürgen Bajorath. Edgeshaper: Bond-centric shapley value-based explanation method for graph  
495 neural networks. *Iscience*, 25(10), 2022.
- 496 [26] Dusit Niyato, Athanasios V Vasilakos, and Zhu Kun. Resource and revenue sharing with  
497 coalition formation of cloud providers: Game theoretic approach. In *2011 11th IEEE/ACM*  
498 *International Symposium on Cluster, Cloud and Grid Computing*, pages 215–224. IEEE, 2011.
- 499 [27] Ki Nohyun, Hoyong Choi, and Hye Won Chung. Data valuation without training of a model. In  
500 *The Eleventh International Conference on Learning Representations*, 2022.
- 501 [28] Jian Pei. A survey on data pricing: from economics to data science. *IEEE Transactions on*  
502 *knowledge and Data Engineering*, 34(10):4586–4608, 2020.
- 503 [29] Stephanie Schoch, Haifeng Xu, and Yangfeng Ji. Cs-shapley: Class-wise shapley values for  
504 data valuation in classification. *Advances in Neural Information Processing Systems*, 35:34574–  
505 34585, 2022.
- 506 [30] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-  
507 Rad. Collective classification in network data. *AI magazine*, 29(3):93–93, 2008.
- 508 [31] Behrooz Shahsavari and Pieter Abbeel. Short-term traffic forecasting: Modeling and learning  
509 spatio-temporal relations in transportation networks using graph neural networks. *University of*  
510 *California at Berkeley, Technical Report No. UCB/EECS-2015-243*, 2015.
- 511 [32] Lloyd S Shapley et al. A value for n-person games. 1953.
- 512 [33] Oleksandr Shchur, Maximilian Mumme, Aleksandar Bojchevski, and Stephan Günnemann.  
513 Pitfalls of graph neural network evaluation. *arXiv preprint arXiv:1811.05868*, 2018.
- 514 [34] Rachael Hwee Ling Sim, Xinyi Xu, and Bryan Kian Hsiang Low. Data valuation in machine  
515 learning: “ingredients”, strategies, and open challenges. In *Proc. IJCAI*, pages 5607–5614, 2022.
- 516 [35] Rafaël Van Belle, Charles Van Damme, Hendrik Tytgat, and Jochen De Weerd. Inductive graph  
517 representation learning for fraud detection. *Expert Systems with Applications*, 193:116463,  
518 2022.
- 519 [36] Vítor V Vasconcelos, Phillip M Hannam, Simon A Levin, and Jorge M Pacheco. Coalition-  
520 structured governance improves cooperation to provide public goods. *Scientific reports*,  
521 10(1):9194, 2020.

- 522 [37] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua  
523 Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- 524 [38] John Von Neumann and Oskar Morgenstern. *Theory of games and economic behavior (60th*  
525 *Anniversary Commemorative Edition)*. Princeton university press, 2007.
- 526 [39] Jiachen T Wang and Ruoxi Jia. Data banzhaf: A robust data valuation framework for machine  
527 learning. In *International Conference on Artificial Intelligence and Statistics*, pages 6388–6421.  
528 PMLR, 2023.
- 529 [40] Eyal Winter. A value for cooperative games with levels structure of cooperation. *International*  
530 *Journal of Game Theory*, 18:227–240, 1989.
- 531 [41] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger.  
532 Simplifying graph convolutional networks. In *ICML*. PMLR, 2019.
- 533 [42] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural  
534 networks? *arXiv preprint arXiv:1810.00826*, 2018.
- 535 [43] Tom Yan and Ariel D Procaccia. If you like shapley then you’ll love the core. In *Proceedings of*  
536 *the AAAI Conference on Artificial Intelligence*, volume 35, pages 5751–5759, 2021.
- 537 [44] Ziyuan Ye, Rihan Huang, Qilin Wu, and Quanying Liu. Same: Uncovering gnn black box  
538 with structure-aware shapley-based multipiece explanations. *Advances in Neural Information*  
539 *Processing Systems*, 36, 2024.
- 540 [45] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure  
541 Leskovec. Graph convolutional neural networks for web-scale recommender systems. In  
542 *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery &*  
543 *data mining*, pages 974–983, 2018.
- 544 [46] Jinsung Yoon, Sercan Arik, and Tomas Pfister. Data valuation using reinforcement learning. In  
545 *International Conference on Machine Learning*, pages 10842–10851. PMLR, 2020.
- 546 [47] Hao Yuan, Haiyang Yu, Jie Wang, Kang Li, and Shuiwang Ji. On explainability of graph neural  
547 networks via subgraph explorations. In *International conference on machine learning*, pages  
548 12241–12252. PMLR, 2021.
- 549 [48] Pu Yuan, Yong Xiao, Guoan Bi, and Liren Zhang. Toward cooperation by carrier aggregation  
550 in heterogeneous networks: A hierarchical game approach. *IEEE Transactions on Vehicular*  
551 *Technology*, 66(2):1670–1683, 2016.
- 552 [49] Shichang Zhang, Yozen Liu, Yizhou Sun, and Neil Shah. Graph-less neural networks: Teaching  
553 old mlps new tricks via distillation. *arXiv preprint arXiv:2110.08727*, 2021.
- 554 [50] Xin Zheng, Yixin Liu, Zhifeng Bao, Meng Fang, Xia Hu, Alan Wee-Chung Liew, and Shirui  
555 Pan. Towards data-centric graph machine learning: Review and outlook. *arXiv preprint*  
556 *arXiv:2309.10979*, 2023.
- 557 [51] Lina Zhou, Shimei Pan, Jianwu Wang, and Athanasios V Vasilakos. Machine learning on big  
558 data: Opportunities and challenges. *Neurocomputing*, 237:350–361, 2017.

## 559 Checklist

560 The checklist follows the references. Please read the checklist guidelines carefully for information on  
561 how to answer these questions. For each question, change the default **[TODO]** to **[Yes]**, **[No]**, or  
562 **[N/A]**. You are strongly encouraged to include a **justification to your answer**, either by referencing  
563 the appropriate section of your paper or providing a brief inline description. For example:

- 564 • Did you include the license to the code and datasets? [Yes] See Section 2.
- 565 • Did you include the license to the code and datasets? [No] The code and the data are
- 566 proprietary.
- 567 • Did you include the license to the code and datasets? [N/A]

568 Please do not modify the questions and only use the provided macros for your answers. Note that the  
 569 Checklist section does not count towards the page limit. In your paper, please delete this instructions  
 570 block and only keep the Checklist section heading above along with the questions/answers below.

571 1. For all authors...

- 572 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s
- 573 contributions and scope? [Yes] The abstract and introduction clearly outline the main
- 574 contributions, including the proposed PC-Winter for graph data valuation and the
- 575 extensive experiments demonstrating its effectiveness.
- 576 (b) Did you describe the limitations of your work? [Yes] The paper mentions that while
- 577 PC-Winter is significantly more efficient than baseline methods like Data Shapley, its
- 578 scalability is still limited and future work on further improving efficiency is desired.
- 579 (c) Did you discuss any potential negative societal impacts of your work? [No] The authors
- 580 believe that the proposed method does not have any potential negative societal impacts.
- 581 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
- 582 them? [Yes]

583 2. If you are including theoretical results...

- 584 (a) Did you state the full set of assumptions of all theoretical results? [Yes] The paper states
- 585 the assumptions and setup for the theoretical results, including the Level Constraint
- 586 and Precedence Constraint in Section 3.2.
- 587 (b) Did you include complete proofs of all theoretical results? [Yes] Complete proofs of the
- 588 Specificity Theorem and Exhaustiveness Theorem regarding permissible permutations
- 589 are provided in Appendix C.

590 3. If you ran experiments (e.g. for benchmarks)...

- 591 (a) Did you include the code, data, and instructions needed to reproduce the main exper-
- 592 imental results (either in the supplemental material or as a URL)? [Yes] The code is
- 593 provided in an anonymous repository (Appendix H.6). The repository also includes
- 594 instructions for downloading the datasets using PyTorch Geometric, allowing for the
- 595 reproduction of the main experimental results.
- 596 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
- 597 were chosen)? [Yes] The datasets and data splits are described in Appendix G.2 and
- 598 G.3. Model hyperparameters and the convergence criteria for the experiments are
- 599 specified in Appendix G.4 and G.5.
- 600 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
- 601 ments multiple times)? [No] Error bars or results from multiple runs are applicable in
- 602 our setting.
- 603 (d) Did you include the total amount of compute and the type of resources used (e.g., type
- 604 of GPUs, internal cluster, or cloud provider)? [Yes] Appendix H.4 includes the total
- 605 GPU hours and hardware used for the experiments.

606 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

- 607 (a) If your work uses existing assets, did you cite the creators? [Yes] The paper cites the
- 608 creators of the benchmark datasets used in the experiments, including Cora, Citeseer,
- 609 Pubmed, Amazon-Photo, Amazon-Computer, and Coauthor-Physics (Appendix G.2).
- 610 (b) Did you mention the license of the assets? [No] The licenses of the datasets are not
- 611 explicitly mentioned.

- 612 (c) Did you include any new assets either in the supplemental material or as a URL? [No]  
613 No new assets are included in the supplemental material or as a URL.
- 614 (d) Did you discuss whether and how consent was obtained from people whose data you're  
615 using/curating? [N/A]
- 616 (e) Did you discuss whether the data you are using/curating contains personally identifiable  
617 information or offensive content? [No] The paper does not discuss if the benchmark  
618 graph datasets used contain personally identifiable information or offensive content.
- 619 5. If you used crowdsourcing or conducted research with human subjects...
- 620 (a) Did you include the full text of instructions given to participants and screenshots, if  
621 applicable? [N/A]
- 622 (b) Did you describe any potential participant risks, with links to Institutional Review  
623 Board (IRB) approvals, if applicable? [N/A]
- 624 (c) Did you include the estimated hourly wage paid to participants and the total amount  
625 spent on participant compensation? [N/A]

## 626 **A Additional Related Work**

627 This section presents an extended review of related works, offering a broader and more nuanced  
628 exploration of the literature surrounding Data Valuation and Graph Neural Networks.

### 629 **A.1 Data Valuation**

630 Data Shapley is proposed in [13] which computes data values with Shapley values in cooperative  
631 game theory. Beta Shapley [20] is a further generalization of Data Shapley by relaxing the efficiency  
632 axiom of the Shapley value. Data Banzhaf [39] offers a data valuation method which is robust to data  
633 noises. Data Valuation with Reinforcement Learning is also explored by [46]. KNN-Shapley [17]  
634 estimates the shapley Value for the K-Nearest Neighbours algorithm in linear time. CS-Shapley [29]  
635 provides a new valuation method that differentiate in-class contribution and out-class contribution.  
636 Data-OOB [21] proposes a data valuation method for a bagging model which leverages the out-of-bag  
637 estimate. Just, Hoang Anh, et al [18] introduce a learning-agnostic data valuation framework by  
638 approximating the utility of a dataset according to its class-wise Wasserstein distance. Another  
639 training-free data valuation method utilizing the complexity-gap score is proposed at the same time  
640 [27]. However, those methods are not designed for the evaluation of data value of graph data which  
641 bears higher complexity due to the interconnections of individual nodes.

### 642 **A.2 Graph Neural Networks**

643 Graph Neural Networks (GNNs) generate informative representations from graph-structured data and  
644 facilitate the solving of many graph-related tasks. Bruna et al. [3] first apply the spectral convolution  
645 operation to graph-structured data. From the spatial perspective, the spectral convolution can be  
646 interpreted to combine the information from its neighbors. GCN [19] simplified this spectral  
647 convolution and proposed to use first-order approximation. Since then, many other attention-based,  
648 sampling-based and simplified GNN variants which follow the same neighborhood aggregation design  
649 have been proposed [37, 14, 12, 41]. Theoretically, those Graph neural networks typically enhance  
650 node representations and model expressiveness through a message-passing mechanism, efficiently  
651 integrating graph data into the learning of representations [42].

### 652 **A.3 Shapley Value in Graph Machine Learning**

653 The Shapley value has found several applications in graph machine learning, primarily in the domain  
654 of explainability for Graph Neural Networks. GraphSVX [8] is one of the early works that utilizes  
655 the Shapley value to explain the predictions of GNNs. It identifies influential nodes and features  
656 for a particular prediction by treating them as players in a cooperative game. However, GraphSVX  
657 focuses on local explanations for individual predictions of a fixed GNN. SubgraphX [47] takes a

658 different approach by explaining GNN predictions through identifying important subgraphs, rather  
 659 than individual nodes or edges. It uses Monte Carlo tree search to efficiently explore different  
 660 subgraphs and proposes to use Shapley values as a measure of subgraph importance. EdgeSHAPer  
 661 [25] is another method that assesses edge importance for GNN predictions using the Shapley value  
 662 concept. It is particularly relevant for molecular graphs where edges represent chemical bonds.  
 663 GNNShap [1] extends upon previous Shapley value based GNN explanation methods by providing  
 664 explanations for edge, leading to better fidelity scores and faster explanations. SAME [44] proposes a  
 665 structure-aware Shapley-based multipiece explanation method for GNNs that can identify important  
 666 substructures and provide explanations composed of multiple connected components.

667 In addition to explainability, Shapley has also been widely adopted for data valuation for conventional  
 668 machine learning methods as discussed in Section 2. However, it has rarely been utilized for data  
 669 valuation on graph data. In this work, we pioneer the exploration of graph data valuation, a challenging  
 670 and previously unexplored problem. Although a recent survey [50] inadvertently refers to GraphSVX  
 671 as a graph data valuation method, it does not align with the traditional definition of data valuation. We  
 672 clarify the key differences between graph data valuation (such as our method) and graph explainability  
 673 (such as GraphSVX) as follows.

- 674 1. In general, data valuation (such as our method) aims to understand how graph elements  
 675 contribute to the model training process, while explainability methods (such as GraphSVX)  
 676 provide post-hoc explanations for a fixed, pre-trained model.
- 677 2. Specifically, our method differs from GraphSVX in several aspects:
  - 678 (a) GraphSVX focuses on the explainability of a local prediction for a single sample, while  
 679 our method aims to quantify the global contribution of graph elements to the overall  
 680 model performance.
  - 681 (b) GraphSVX operates post hoc, analyzing the contributions of features and nodes in the  
 682 *testing graph* to the predictions of an already-trained GNN model, while our approach  
 683 focuses on the global contribution of each data element in the *training graph* to the  
 684 GNN model’s training process.
  - 685 (c) GraphSVX employs the standard Shapley value formulation, which assumes free  
 686 collaboration among players, while our work introduces the PC-Winter value to  
 687 handle the unique hierarchical coalition structures inherent in graph data valuation.

688 To the best of our knowledge, our investigation constitutes the first foray into graph data valuation,  
 689 pioneering research in this previously uncharted domain.

## 690 B Mathematical Formulation of Winter Value

691 The Shapley value offers a solution for equitable payoff distribution in cooperative games, assuming  
 692 that players cooperate without any predefined structure. In reality, however, cooperative games often  
 693 have inherent hierarchical coalitions. To accommodate these structured coalitions, the Winter value  
 694 extends Shapley value to handle this extra coalition constraints.

695 Specifically, considering level structures  $\mathcal{B}$ , with  $\mathcal{B} = B_0, \dots, B_n$  representing a sequence of player  
 696 partitions. Here, a partition,  $B_m$ , subdivides the player set  $\mathcal{P}$  into a set of disjoint, non-empty subsets  
 697  $T_1, T_2, \dots, T_k$ . These disjoint subsets satisfy the condition that their union reconstructs the original  
 698 player set  $\mathcal{P}$ , which means  $T_1 \cup T_2 \cup \dots \cup T_k = \mathcal{P}$ . This partition sequence forms a hierarchy where  
 699  $B_0$  represents individual players as the leaves of the structure and  $B_n$  functions as the root of this  
 700 hierarchy.

701 We then determine  $\Omega(\mathcal{B})$ , the set of all permissible permutations, starting with a single partition  $B_m$ :

$$\Omega(B_m) = \{\pi \in \Pi(\mathcal{P}) : \forall T \in B_m, \forall i, j \in T \text{ and } k \in \mathcal{P}, \\ \text{if } \pi(i) < \pi(k) < \pi(j) \text{ then } k \in T\}.$$

702  $\Omega(\mathcal{B})$  can be further defined as the set of permutations which satisfy all constraints of all levels,  
 703  $\Omega(\mathcal{B}) = \bigcap_{t=0}^n \Omega(B_t)$ .

A permissible permutation  $\pi$  from the set  $\Omega(\mathcal{B})$  requires that players from any derived coalition of  $\mathcal{B}$  must appear consecutively. Given the defined set of permissible permutations  $\Omega(\mathcal{B})$ , the Winter value  $\Phi$  for player  $i$  is calculated as:

$$\Phi_i(\mathcal{P}, U, \mathcal{B}) = \frac{1}{|\Omega(\mathcal{B})|} \sum_{\pi \in \Omega(\mathcal{B})} (U(\mathcal{P}_i^\pi \cup i) - U(\mathcal{P}_i^\pi))$$

704 where  $\mathcal{P}_i^\pi = \{j \in N : \pi(j) < \pi(i)\}$  is the set of predecessors of  $i$  at the permutation  $\sigma$  and  $U$  is the  
705 utility function in the cooperative game.

## 706 C Proofs of Theorems

707 **Theorem 1** (Specificity). *Given a contribution tree  $\mathcal{T}$  with a set of players  $\mathcal{P}$ , any DFS traversal*  
708 *over the  $\mathcal{T}$  results in a permissible permutation of  $\mathcal{P}$  that satisfies both the Level Constraint and*  
709 *Precedence Constraint.*

710 *Proof.* We validate the theorem by demonstrating that a permutation obtained through pre-order  
711 traversal on  $\mathcal{T}$  meets Level Constraints and Precedence Constraints. (1) Level Constraints: During a  
712 pre-order traversal of  $\mathcal{T}$ , a node  $p$  and its descendants  $\mathcal{D}(p)$  are visited sequentially before moving  
713 to another subtree. Thus, in the resulting permutation  $\pi$ , the positions of  $p$  and any  $i, j \in \mathcal{D}(p)$   
714 are inherently close to each other, satisfying the condition  $|\pi[i] - \pi[j]| \leq |\mathcal{D}(p)|$ . This contiguous  
715 traversal ensures that all descendants and the node itself form a continuous sequence in  $\pi$ , meeting  
716 the Level Constraint. (2) Precedence Constraints: In the same traversal, each node  $p$  is visited before  
717 its descendants. Therefore, in  $\pi$ , the position of  $p$  always precedes the positions of its descendants,  
718 i.e.,  $\pi[p] < \pi[i]$  for all  $i \in \mathcal{D}(p)$ . This traversal pattern naturally embeds the hierarchy of the tree  
719 into the permutation, ensuring that ancestors are positioned before their descendants, in line with the  
720 Precedence Constraint.  $\square$

721 **Theorem 2** (Exhaustiveness). *Given a contribution tree  $\mathcal{T}$  with a set of players  $\mathcal{P}$ , any permissible*  
722 *permutation  $\pi \in \Omega$  can be generated by a corresponding DFS traversal of  $\mathcal{T}$ .*

723 *Proof.* To prove the theorem of exhaustiveness, consider a contribution tree  $\mathcal{T}$  with a set of players  
724  $\mathcal{P}$  and any permissible permutation  $\pi \in \Omega$ . We apply induction on the depth of  $\mathcal{T}$ . For the base case,  
725 when  $\mathcal{T}$  has a depth of 1, which means there are no dependencies among players, any permissible  
726 permutation of players is trivially generated by a DFS traversal since there are no constraints on the  
727 order of traversal. For the inductive step, assume the theorem holds for contribution trees of depth  $k$ .  
728 For a contribution tree of depth  $k + 1$   $\mathcal{T}^{k+1}$ , consider its root node and subtrees of depth  $k$  rooted  
729 at the child nodes of the root node. For any given permissible permutation  $\pi$  corresponding to the  
730  $\mathcal{T}^{k+1}$ , according to the Level Constraint, it is a direct composition of the permissible permutations  
731 corresponding to the subtrees of depth  $k$  rooted at the child nodes of the root node. Now we can  
732 construct a DFS traversal over the contribution tree  $\mathcal{T}^{k+1}$  that can generate  $\pi$ . Specifically, the order  
733 of composition defines the traversal order of the child nodes of the root node. Furthermore, by the  
734 inductive hypothesis, any permissible permutations corresponding to the subtrees can be generated  
735 by DFS traversal over the subtrees. Hence, at each child node of the root node, we just follow the  
736 corresponding DFS traversal of its corresponding tree. This DFS traversal can generate the given  
737 permutation  $\pi$ , which completes the proof.  $\square$

## 738 D Hierarchical Truncation

739 In Table 1, we present data comparing the number of model re-trainings on the all six dataset with  
740 and without the application of truncation. For the Citeseer dataset, the truncation ratios are defined as  
741 1st-hop: 0.5 and 2nd-hop: 0.7. For the remaining datasets, the truncation ratios are set at 1st-hop: 0.7  
742 and 2nd-hop: 0.9. The results clearly indicate that the number of model re-trainings is substantially  
743 reduced when truncation is applied. For instance, focusing on the Citeseer dataset the application  
744 of truncation significantly reduces the number of retrainings from 1388 to 535. This significant



745 decrease, especially in larger datasets like Amazon-Photo and Amazon-Computer, where retraining  
 746 instances decrease from 147664 to 6258 and from 317959 to 12139 respectively, can be attributed to  
 747 the substantial number of 2-distance neighbors present in these datasets. The application of truncation  
 748 effectively reduces the computation by omitting a considerable portion of these neighbors. This  
 749 finding also implies that overall training time is decreased while still maintaining the ability to  
 750 accurately measure the total marginal contribution.

Table 1: Retraining Number Comparison Per Permutation

Dataset	w.o. Truncation	w.t. Truncation
Cora	2241	756
Citeseer	1388	535
Pubmed	3683	887
Amazon-Photo	147664	6258
Amazon-Computer	317959	12139
Coauthor-Physics	11178	852

## 751 E Mixed Node Dropping Experiment

752 As mentioned in the experiment, labeled nodes will dominate the performance curve when both  
 753 labeled nodes and unlabeled nodes. The corresponding experiment result is shown in the Figure 6.  
 754 This experiment validates the assumption that a effective data valuation method would naturally rank  
 755 labeled nodes for earlier removal over their unlabeled counterparts. For instance, in the Cora dataset,  
 756 we can observe that the initial drop in accuracy is significant, indicating the removal of high-value  
 757 labeled nodes. As the experiment progresses and more nodes are removed, the accuracy barely  
 758 changes, reflecting the removal of unlabeled nodes which has a minimal impact on performance  
 759 when most labeled nodes are unavailable. The observed pattern across all datasets is consistent: there  
 760 is a substantial drop in performance at the beginning, followed by a plateau with minimal changes.  
 761 This suggests that the initial set of nodes removed, predominantly high-value labeled nodes, are  
 762 those critical to the model’s performance, whereas the subsequent nodes show less influence on the  
 outcome.




Figure 6: Mixed Node Dropping Experiment

763

764 **F Labeled Node Dropping Experiment**

765 Here, we perform node dropping experiment employing the aggregated value define in the main-  
 766 body of paper, to demonstrate that PC-Winter can capture the heterogeneous influence of labeled  
 767 nodes. As shown in the Figure 7, both PC-Winter and Data Shapley demonstrate effectiveness  
 768 in capturing the diverse contributions of labeled nodes to the model’s performance. Particularly in  
 769 the Pubmed and Amazon-Photo datasets, PC-Winter exhibits better performance compared to Data  
 770 Shapley. In other datasets, such as Cora, Citeseer, and Coauthor-Physics, PC-Winter shows results  
 771 that are on par with Data Shapley.

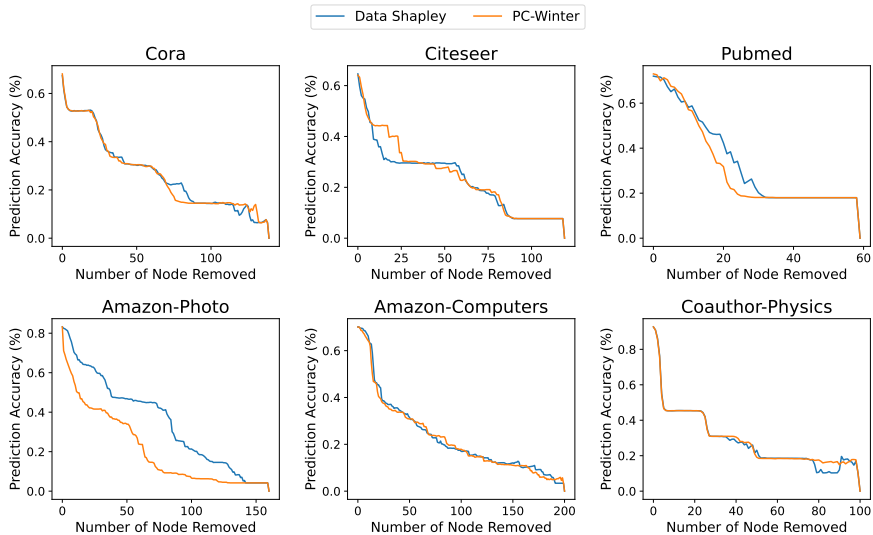


Figure 7: Labeled Node Dropping Experiment

772 **G Experimental Details**

773 **G.1 Inductive Setting**

774 Our experiments focus on the inductive node classification task, which aims to generalize a trained  
 775 model to unseen nodes and is commonly adopted in real-world graph applications [14, 35, 16, 9].  
 776 Unlike the transductive setting [19] which incorporates the test nodes in the model training process,  
 777 the inductive setting separates them apart from the training graph. Such a separation allows us to  
 778 measure the value of the graph elements in the training graph solely based on their contribution  
 779 to GNN model training. Following [14], we split each graph  $\mathcal{G}$  into 3 disjoint subgraphs: training  
 780 graph  $\mathcal{G}_{tr}$ , validation graph  $\mathcal{G}_{va}$ , and test graph  $\mathcal{G}_{te}$ . The training graph  $\mathcal{G}_{tr}$  is constructed without  
 781 any nodes from the validation or test set. Correspondingly, edges connecting to a validation node  
 782 or a testing node are also removed from the training graph. For the validation graph  $\mathcal{V}_{va}$  and the  
 783 testing graph  $\mathcal{V}_{te}$ , only edges with both nodes within the respective node sets are retained, which is  
 784 aligned with the inductive setting in prior work [49]. We utilize  $\mathcal{G}_{tr}$  to train the GNN model, which is  
 785 evaluated on  $\mathcal{V}_{va}$  for obtaining the data values for elements. The test graph  $\mathcal{V}_{te}$  is utilized to evaluate  
 786 the effectiveness of the obtained values.

787 **G.2 Datasets**

788 We assess the proposed approach on six real-world benchmark datasets. These include three citation  
 789 graphs, Cora, Citeseer and Pubmed [30] and two Amazon Datasets, Amazon-Photo and Amazon-  
 790 Computer, and Coauthor-Physics [33]. The detailed statistics of datasets are summarized in Table  
 791 2.

Table 2: Dataset Summary

Dataset	# Node	# Edge	# Class	# Feature	# Train/Val/Test
Cora	2,708	5,429	7	1,433	140 / 500 / 1,000
Citeseer	3,327	4,732	6	3,703	120 / 500 / 1,000
Pubmed	19,717	44,338	3	500	60 / 500 / 1,000
Amazon-Photo	7,650	119,081	8	745	160 / 20% / 20%
Amazon-Computer	13,752	245,861	10	767	200 / 20% / 20%
Coauthor-Physics	34,493	247,962	8	745	100 / 20% / 20%

### 792 G.3 Dataset Split

793 In the conducted experiments, we split each graph  $\mathcal{G}$  into 3 disjoint subgraphs: training graph  $\mathcal{G}_{tr}$ ,  
 794 validation graph  $\mathcal{G}_{va}$ , and test graph  $\mathcal{G}_{te}$ . The training graph  $\mathcal{G}_{tr}$  is constructed without any nodes  
 795 from the validation or test set. Correspondingly, edges connecting to a validation node or a testing  
 796 node are also removed from the training graph. For the validation graph  $\mathcal{V}_{va}$  and the testing graph  
 797  $\mathcal{V}_{te}$ , only edges with both nodes within the respective node sets are retained, which is aligned with the  
 798 inductive setting in prior work [49]. We utilize  $\mathcal{G}_{tr}$  to train the GNN model, which is evaluated on  $\mathcal{V}_{va}$   
 799 for obtaining the data values for elements. The test graph  $\mathcal{V}_{te}$  is utilized to evaluate the effectiveness  
 800 of the obtained values. In the case of the specific split for each dataset, for the citation networks, we  
 801 adopt public train/val/test splits in our experiments. For the remaining datasets, we randomly select  
 802 20 labeled nodes per class for training, 20% nodes for validation and 20% nodes as the testing set.

### 803 G.4 Convergence Criteria

**Convergence Criterion.** For permutation-based data valuation methods such as Data Shapley and PC-Winter, we follow convergence criteria similar to the one applied in prior work [13] to determine the number of permutations for approximating data values:

$$\frac{1}{n} \sum_{i=1}^n \frac{|v_i^t - v_i^{t-20}|}{|v_i^t|} < 0.05$$

804 where  $v_i^t$  is the estimated value for the data element  $i$  using the first  $t$  sampled permutations.

805 **Time Limit.** For larger datasets, sampling a sufficient number of permutations for converged data  
 806 values could be impractical in time. To address this and to stay within a realistic scope, we cap the  
 807 computation time at 120 GPU hours on NVIDIA Titan RTX, after which the calculation is terminated.

### 808 G.5 Truncation Ratios and Hyper-parameters

809 Table 3 includes the hyper-parameters and truncation ratios used for value estimation.

Table 3: Truncation Ratios and Hyper-parameters

Dataset	Truncation Ratio	Learning Rate	Epoch	Weight Decay
Cora	0.5-0.7	0.01	200	5e-4
Citeseer	0.5-0.7	0.01	200	5e-4
Pubmed	0.5-0.7	0.01	200	5e-4
Amazon-Photo	0.7-0.9	0.1	200	0
Amazon-Computer	0.7-0.9	0.1	200	0
Coauthor-Physics	0.7-0.9	0.01	30	5e-4

### 810 G.6 Baselines

#### 811 G.6.1 Dropping High-Value Nodes

812 Here, we introduce the baselines used for comparison to validate the effectiveness of the proposed  
 813 method in the dropping node experiment:

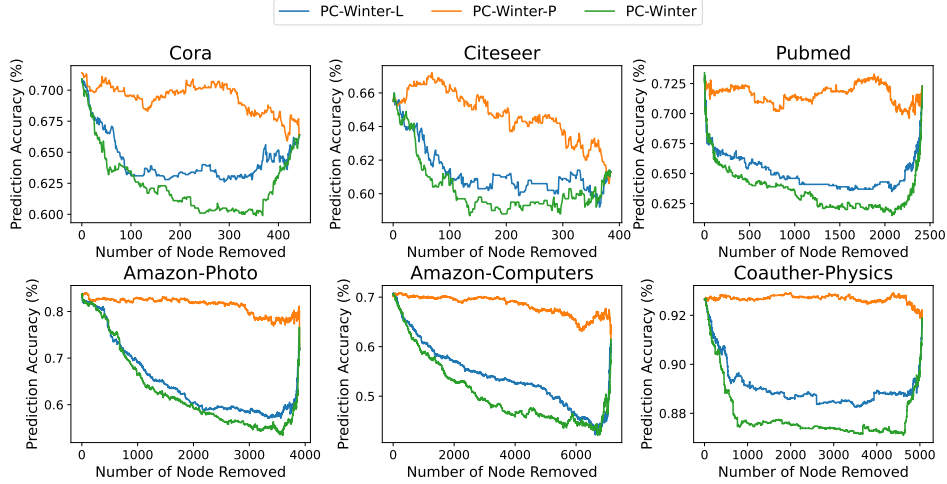


Figure 8: Ablation Study

- 814 • **Random Value:** It assigns nodes with random values, which leads to random ranking without any  
815 specific pattern or correlation to the node’s features.
  - 816 • **Degree-based Value:** A node is assigned its degree as its value, assuming that a node’s impor-  
817 tance in the graph is indicated by its degree.
  - 818 • **Leave-one-out (L00):** This method calculates a node’s value based on its marginal contribution  
819 compared to the rest of the training nodes. Specifically, the value  $v(i)$  assigned to each node  $i$   
820 is its marginal utility, calculated as  $v(i) = U(\mathcal{G}_{tr}) - U(\mathcal{G}_{tr}^{-i})$ , where  $\mathcal{G}_{tr}^{-i}$  denotes the training  
821 graph excluding node  $i$ . The utility function  $U$  measures the model’s validation performance when  
822 trained on the given graph. In essence, the drop in performance due to the removal of a node is  
823 treated as the value of that node.
  - 824 • **Data Shapley:** The node values are approximated with the Monte Carlo sampling method of  
825 Data Shapley [13] by treating both labeled nodes and unlabeled nodes as players. Notably, we  
826 only include those unlabeled nodes within the 2-hop neighbors of labeled nodes in the evaluation  
827 process. There are two approximation methods: Truncated Monte Carlo approximation and  
828 Gradient Shapley in [13]. We adopt the Truncated Monte Carlo approximation as it consistently  
829 outperforms the other variants in various experiments.
- 830 Notably, there is a recent work [6] that aims at characterizing the impact of elements on model  
831 performance. Their goal is to approximate L00 value. Thus, we do not include it as a baseline as L00  
832 is already included.

### 833 G.6.2 Adding High-Value Edges

834 Here are the detailed descriptions on the baselines applied in the edge adding experiment.

- 835 • **Random Value:** it assigns edges with random values, reflecting a baseline where no information  
836 are used for differentiating the importance of edges.
- 837 • **Edge-Betweenness:** the Edge-Betweenness of an edge  $e$  is the the fraction of all pairwise shortest  
838 paths that go through  $e$ . This classic approach assesses an edge’s importance based on its role in  
839 the overall network connectivity.
- 840 • **Leave-one-out (L00):** This method calculates a edge  $e$ ’s value  $v(e)$  based on its marginal  
841 contribution compared to the rest of the training graph. In specific,  $v(e) = U(\mathcal{G}_{tr}) - U(\mathcal{G}_{tr}^{-e})$   
842 Here,  $e \in \mathcal{G}_{tr}$  represents an edge in the training graph  $\mathcal{G}_{tr}$ , and  $\mathcal{G}_{tr}^{-e}$  refers to the training graph  
843 excluding the edge  $e$ .

## 844 H Ablation Study and Parameter Analysis

### 845 H.1 Ablation Study

846 This Appendix Section offers an in-depth ablation analysis across full six datasets to investigate  
847 the necessity of both Level Constraint and Precedence Constraint in defining an effective graph  
848 value. The results, as shown in Figure 8, consistently demonstrate across all datasets that the absence  
849 of either constraint leads to a degraded result when compared to the one incorporating both. This  
850 underscores the importance of both two constraints in capturing the contributions of graph elements  
851 to overall model performance.

### 852 H.2 The Impact of Permutation Number

853 This part expands upon the permutation analysis presented in the main paper. It provides comprehen-  
854 sive results across various datasets, illustrating how different numbers of sample permutations impact  
the accuracy of PC-Winter. The results of full datasets are shown in Figure 9. The results reveals

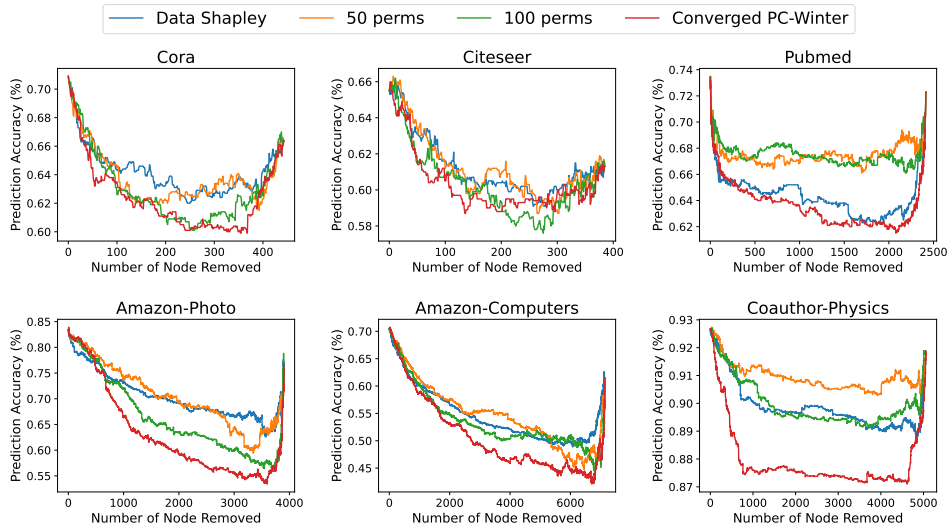


Figure 9: The Impact of Permutation Numbers

855 that increasing the number of permutations generally improves the performance and accuracy of the  
856 valuation. PC-Winter also show robust results even with a limited number of permutations, high-  
857 lighting its effectiveness. The phenomenon is consistent across all datasets where our approach with  
858 just 50 to 100 permutations manages to compete closely with the fully converged Data Shapley,  
859 emphasizing the efficiency of PC-Winter in various settings.  
860

### 861 H.3 The Impact of Truncation Ratios

862 Our approach involves truncating the iterations involving the first and second-hop neighbors of a  
863 labeled node during value estimation. Here, we investigate the impact of truncation proportion  
864 on overall performance, using the same number of permutations as in our primary node-dropping  
865 experiment. As shown in Figure 10, we adjusted the truncation ratios for the Citation Network  
866 datasets. The ratios ranged from truncating 50% of the first-hop and 70% of the second-hop neighbors  
867 (0.5-0.7), up to 90% truncation for either first-hop (0.9-0.7) or second-hop (0.5-0.9) neighbors. For  
868 the Cora and Citeseer datasets, increasing truncation at the first-hop level had a minimal impact  
869 on performance, and PC-Winter still significantly outperformed Data Shapley. In the case of  
870 the Pubmed dataset, more extensive truncation at the first-hop level notably reduced performance.  
871 Regarding large datasets such as the Amazon, while truncation at either the first or second-hop  
872 levels had a marginal negative effect on performance, PC-Winter’s estimated data values generally  
873 remained superior to results of Data Shapley.

874 In addition, we provide a detailed analysis of our truncation strategy across other datasets. It includes  
 875 results not presented in the main text, focusing on the impact of limiting model retraining times to the  
 876 first and second-hop neighbors in value estimation. We investigate the impact of truncation proportion  
 877 on overall performance, using the same number of permutations as in our primary node-dropping  
 878 experiment. The findings on full datasets are illustrated in Figure 10. Specifically, our findings reveal  
 879 that in datasets like Cora and Citeseer, adjusting truncation primarily at the first-hop level has a  
 880 negligible impact on the accuracy of node valuation, with PC-Winter still maintaining a considerable  
 881 advantage over Data Shapley. For large datasets such as the Amazon-Photo, Amazon-Computers  
 882 and Coauthor-Physics, while truncations had a marginal negative effect on performance, PC-Winter  
 883’s estimated data values generally remained better than Data Shapley. This analysis indicates that  
 884 PC-Winter can afford to employ larger truncation, enhancing computational efficiency without  
 885 substantially sacrificing the quality of data valuation.

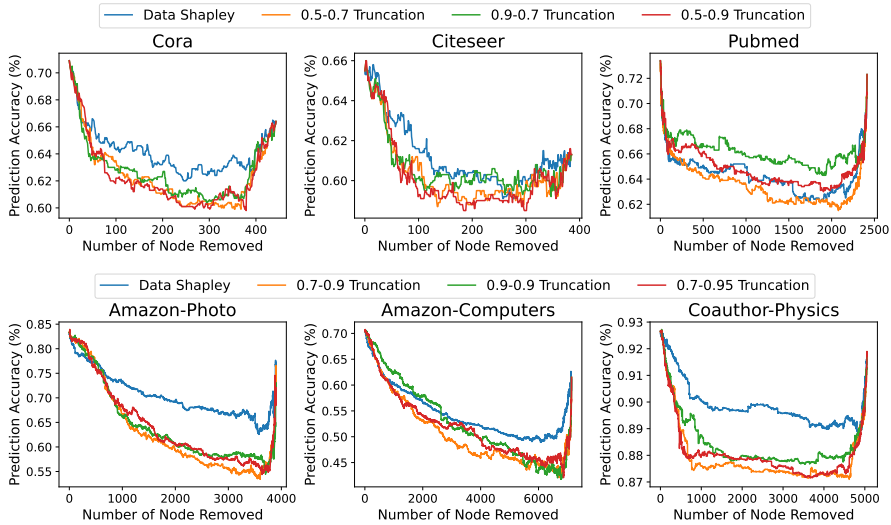


Figure 10: The Impact of Truncation Ratios

Table 4: Permutation Number and Time Comparison

Dataset	Truncation	PC-Winter		Data Shapley	
		Perm Number	Perm Time (hrs)	Perm Number	Perm Time (hrs)
Cora	0.5-0.7	325	0.013	327	0.024
Citeseer	0.5-0.7	291	0.018	279	0.037
Pubmed	0.5-0.7	316	0.025	281	0.285
Amazon-Photo	0.7-0.9	418	0.211	109	1.105
Amazon-Computer	0.7-0.9	181	0.662	33	3.566
Coauthor-Physics	0.7-0.9	460	0.119	45	2.642

886 **H.4 Efficiency Analysis**

887 Here, we compare the computational efficiency of our proposed method PC-Winter and the Data  
 888 Shapley approach in terms of permutation number and time per permutation. As detailed in Table 4,  
 889 the results indicate that PC-Winter requires significantly less time to compute each permutation  
 890 across various datasets. Specifically, for the Cora dataset, PC-Winter completes each permutation  
 891 in approximately half the time required by Data Shapley. Moving to larger datasets, the  
 892 efficiency of PC-Winter becomes even more pronounced. For instance, in the Amazon-Computer  
 893 dataset, PC-Winter’s permutation time is only a fraction of what is required by Data Shapley  
 894 —PC-Winter takes slightly over half an hour per permutation whereas Data Shapley exceeds three

895 and a half hours. This consistent reduction in permutation time demonstrates the computational  
896 advantage of PC-Winter, particularly when handling large graphs. Combining the insights from the  
897 Permutation Analysis shown in Figure 9 with the Permutation Comparison Table 4, we observe that  
898 for datasets such as Cora, Citeseer, Amazon-Photo, and Amazon-Computer, around 50 permutations  
899 are sufficient for PC-Winter to achieve performance comparable to that of Data Shapley. Simple  
900 calculations demonstrate that our method is significantly faster than Data Shapley in achieving  
901 similar performance levels. For instance, in the Cora dataset, the speedup factor is  $\frac{327 \times 0.024}{50 \times 0.013} = 12.07$ ,  
902 and for the Citeseer dataset, it is  $\frac{279 \times 0.037}{50 \times 0.018} = 11.47$ . The speedup factors for Amazon-Photo and  
903 Amazon-Computer are  $\frac{109 \times 1.105}{50 \times 0.211} = 11.42$ , and  $\frac{33 \times 3.566}{50 \times 0.662} = 3.57$ , respectively. For Coauthor-Physics,  
904 it takes about 100 permutations for PC-Winter to match the performance of Data Shapley, which  
905 implies a speedup factor of  $\frac{45 \times 2.642}{100 \times 0.119} = 10.00$ . In conclusion, PC-Winter can achieve stronger  
906 performance than Data Shapley using the same or even less time. Furthermore, it takes PC-Winter  
907 much less time to achieve comparable performance as Data Shapley. Notably, though PC-Winter  
908 is significantly more efficient than Data Shapley, its scalability is still limited, and future work in  
909 further improving its efficiency is desired.

## 910 H.5 Complexity Analysis

911 We analyze the complexity of the PC-Winter. For convenience, we assume that we are dealing with a  
912  $d$ -regular graph. There are a total of  $L$  labeled nodes in the graph. As described in the paper, we deal  
913 with a GNN model with 2 layers. Without loss of generality, we use  $F$  to denote the dimensionality  
914 of node representations in each layer. We assume the number of classes in the dataset is  $C$ . For  
915 hierarchical truncation, we assume we adopt a truncation ratio of  $r_1 - r_2$ , which is consistent with  
916 the description in Section 3.4. Then, the number of nodes in a computation tree for any labeled  
917 node is  $N_{full} = 1 + d + d^2$ . With hierarchical truncation, the number of nodes in the truncated  
918 computation tree is  $N_{trun} = 1 + d \cdot (1 - r_1) + d^2 \cdot (1 - r_1)(1 - r_2)$ . When the truncation ratios  
919 are large,  $N_{trun} \ll N_{full}$ . For instance, when  $r_1 = r_2 = 0.9$ ,  $N_{trun}$  could be less than 5Time  
920 Complexity Analysis: We now analyze the time complexity of a single permissible permutation of the  
921 PC-Winter algorithm. We begin by examining the time complexity of generating a single permissive  
922 permutation. Then, we investigate the complexity of a single model retraining and provide the total  
923 retraining number for a single permutation. Finally, we combine these analyses to derive the overall  
924 time complexity for generating one permissible permutation and going through it for calculating the  
925 marginal contributions. Time complexity of generating a single permissive permutation: The time  
926 complexity of traversing the truncated contribution tree to generate a single permissive permutation is  
927  $O(L \cdot N_{trun})$ . In particular, there are  $L \cdot N_{trun} + 1$  nodes in the contribution tree (including the dummy  
928 node). Hence, the cost of a DFS traversal over the contribution tree is  $O(L \cdot N_{trun} + 1 + L \cdot N_{trun})$   
929  $= O(L \cdot N_{trun})$ . Time complexity of one model retraining: As described in Section 3.4, with local  
930 propagation, for each model retraining, we only need to perform feature aggregation on a single partial  
931 computation tree. The size of a partial computation tree is, on average,  $\frac{N_{trun}}{2}$ . Therefore, the feature  
932 aggregation complexity for each retraining step is  $O(\frac{N_{trun}}{2} \cdot F)$ , where  $F$  is the dimension of node  
933 features. The feature transformation complexity for each model retraining is  $O(F \cdot F + F \cdot C) = O(F^2)$ ,  
934 where  $C$  is the output dimension (number of classes) of the GNN model. Therefore, the total time  
935 complexity of a single retraining is  $O(\frac{N_{trun}}{2} \cdot F + F^2)$ . Without local propagation, the feature  
936 aggregation complexity for each model retraining would be much larger, since the propagation needs  
937 to be performed on the entire graph. The number of model retraining in a single permutation: In a  
938 permissible permutation, we need to perform retraining for each node in the truncated contribution  
939 tree, which has  $L \cdot N_{trun}$  nodes in total. Therefore,  $L \cdot N_{trun}$  model retrainings are needed for a single  
940 permutation. Total time complexity for a single permissible permutation: With local propagation and  
941 hierarchical truncation, the total time complexity of a single permissible permutation in PC-Winter is:  
942  $O(L \cdot N_{trun} + L \cdot N_{trun} \cdot (\frac{N_{trun}}{2} \cdot F + F^2)) = O(L \cdot N_{trun} \cdot (1 + \frac{N_{trun}}{2} \cdot F + F^2)) = O(L \cdot N_{trun} \cdot (\frac{N_{trun}}{2} \cdot F + F^2))$ .  
943 Notably, the time complexity of generating a permissible permutation is negligible compared to the  
944 cost of model retraining. The proposed strategies, hierarchical truncation, and local propagation, help  
945 reduce the overall time complexity of the PC-Winter algorithm. In particular, hierarchical truncation  
946 makes  $N_{trun}$  much smaller than  $N_{full}$ , greatly decreasing the total number of model retraining required

947 for a single permutation. On the other hand, Local propagation reduces the feature aggregation  
948 complexity, greatly reducing the cost of each retraining.

#### 949 **H.6 Code Availability**

950 To facilitate the reproducibility of our work and to encourage further research in the field of graph  
951 data valuation, we have made our code publicly available on an anonymous repository at <https://anonymous.4open.science/r/graph-data-valuation-B348>. The repository contains the  
952 implementation of the PC-Winter algorithm, along with scripts for running the experiments presented  
953 in this paper. We welcome researchers and practitioners to utilize and build upon our code for their  
954 own research and applications in graph data valuation.  
955