

---

# Reinforcement Learning Enhanced Explainer for Graph Neural Networks

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 Graph neural networks (GNNs) have recently emerged as revolutionary technologies for machine learning tasks on graphs. In GNNs, the graph structure is generally  
2 incorporated with node representation via the message passing scheme, making the  
3 explanation much more challenging. Given a trained GNN model, a GNN explainer  
4 aims to identify a most influential subgraph to interpret the prediction of an instance  
5 (e.g., a node or a graph), which is essentially a combinatorial optimization problem  
6 over graph. The existing works solve this problem by continuous relaxation  
7 or search-based heuristics. But they suffer from key issues such as violation of  
8 message passing and hand-crafted heuristics, leading to inferior interpretability. To  
9 address these issues, we propose a RL-enhanced GNN explainer, *RG-Explainer*,  
10 which consists of three main components: starting point selection, iterative graph  
11 generation and stopping criteria learning. RG-Explainer could construct a connected  
12 explanatory subgraph by sequentially adding nodes from the boundary of  
13 the current generated graph, which is consistent with the message passing scheme.  
14 Further, we design an effective seed locator to select the starting point, and learn  
15 stopping criteria to generate superior explanations. Extensive experiments on both  
16 synthetic and real datasets show that RG-Explainer outperforms state-of-the-art  
17 GNN explainers. Moreover, RG-Explainer can be applied in the inductive setting,  
18 demonstrating its better generalization ability.  
19

## 20 1 Introduction

21 Graph Neural Networks (GNNs) extend neural network models on ubiquitous graph data via utilizing  
22 the message passing scheme to incorporate graph structures with node features. They have achieved  
23 state-of-the-art performance not only in classic machine learning tasks on graphs, e.g., node classification  
24 [10, 25], link prediction [35], and graph classification [30], but also in reasoning tasks, e.g.,  
25 intuitive physics [4], mathematical reasoning [23], and IQ tests [3]. Similar to most deep learning  
26 methods, one major limitation of GNNs is the lack of the interpretability for the predicted results; a  
27 post-hoc analysis is usually needed to explain the results.

28 To enhance the interpretability of GNNs, a line of works [31, 16, 27, 34, 28] focused on developing  
29 GNN explainers. The goal of GNN explainers is to identify a most influential subgraph structure to  
30 interpret the predicted label of an instance (e.g., a node or a graph). It can be generally formulated as  
31 an optimization problem that maximizes the mutual information between the predicted results and  
32 the distribution of relevant subgraphs under some size constraints.

33 The pioneering works, e.g., GNNExplainer [31] and PGExplainer [16], attempt to solve the optimization  
34 problem with continuous relaxation. These methods optimize a soft mask matrix for edges,  
35 and select the important nodes/edges by the threshold. However, they cannot guarantee that nodes  
36 and edges in the output subgraph are connected. Thus, their explanatory subgraphs cannot explicitly

37 visualize the message passing paths. Besides, they consider the importance of each edge indepen-  
38 dently, and ignore the interactions among selected nodes and edges. Some recent works, such as  
39 SubgraphX [34] and Causal Screening [28], design the search criteria and use search-based methods  
40 to solve the optimization problem. Due to the combinatorial property of searching explanatory graph  
41 structures, it is difficult to design a general hand-crafted search criterion. These criteria are limited on  
42 specific situations and thus not widely applicable.

43 To address these issues, we propose RG-Explainer, which adopts reinforcement learning to explain  
44 GNNs’ predictions. Our framework is inspired by classic combinatorial optimization solvers, which  
45 consists of three crucial steps: *starting point selection*, *iterative graph generation* and *stopping*  
46 *criteria learning*. These three components work together to generate an explanatory graph that  
47 interprets the predicted label of a given node/graph instance, as we elaborate next.

48 Firstly, starting point selection selects the most important node as the seed node in the instance. If the  
49 task is to interpret the prediction of a specific node label, then the most important node refers to the  
50 node itself. To explain a graph label, we design a seed locator to learn the node that influences the  
51 graph label the most. Iterative graph generation is the key module in our method, which generates the  
52 nodes in the explanatory graph sequentially. Specifically, we add an influential node (action) from  
53 the neighbors based on the current generated graph (state) at each step. It explicitly guarantees the  
54 connectivity of the generated graph. The generation process is controlled by the reward, i.e., the  
55 mutual information between the original predicted label and the label made by the generated graph.  
56 To ensure a compact and meaningful explanatory graph, we also involve some constraints into the  
57 reward, such as size loss, radius penalty and similarity loss. Finally, stopping criteria are learned to  
58 further avoid generating very large explanatory graphs.

59 Furthermore, our method has better generalization ability and can be applied in both transductive  
60 and inductive setting. Different from the search-based methods, we learn the heuristics from the data  
61 automatically. A well-trained RG-Explainer can infer the explanations of instances which are not  
62 involved in the training phase.

63 We conduct extensive experiments on both synthetic and real-world datasets to show that the proposed  
64 RG-Explainer can achieve superior performance compared to state-of-the-art GNN explainers. In  
65 particular, our visualization results further demonstrate the better interpretability of our method.

## 66 2 Related Work

67 **Graph Neural Networks.** Graph neural networks (GNNs) have achieved great success on non-  
68 Euclidean data, e.g., graphs. The majority of GNNs used today follow the message passing  
69 scheme [8], which aggregate information from neighbors with different aggregation functions,  
70 like mean/max/LSTM-pooling in GCN [12] and GraphSAGE [10], sum-pooling in GIN [30], at-  
71 tention mechanisms in GAT [25], etc. SGC [29] observes that the superior performance of GNNs  
72 is mainly due to the neighbor aggregation rather than feature transformation and nonlinearity, and  
73 proposed a simple and fast GNN model. APPNP [13] shares the similar idea by decoupling feature  
74 transformation and neighbor aggregation.

75 **Graph Generation.** There have been a variety of methods for graph generation. RVAE [19]  
76 is a variational auto-encoder (VAE) based method with a regularizer to ensure semantic validity.  
77 Normalizing flow based methods including GraphNVP [20], GraphAF [24], and GraphDF [17] utilize  
78 invertible neural networks to define mappings between latent variables and data points. Generative  
79 adversarial networks (GANs) [9] based methods like MolGAN [7] and GCPN [32] involve a generator  
80 and a discriminator, where the generator is adversarially trained to fool the discriminator.

81 From the perspective of graph generation process, they can be classified into one-shot generation  
82 and iterative generation. RVAE and MolGAN directly generate adjacency matrices, while GraphAF,  
83 GraphDF and GCPN generate graphs by sequentially adding new nodes and edges. Though our  
84 proposed RG-Explainer is an iterative generation method, RG-Explainer is different from the graph  
85 generation methods in that the above methods generate graph out of the air, while RG-Explainer  
86 needs to dynamically select suitable subgraphs to explain the predictions.

87 **Graph Combinatorial Optimization with RL.** With the success of deep reinforcement learning  
88 in games [21], researchers have attempted to utilize RL techniques for the graph combinatorial

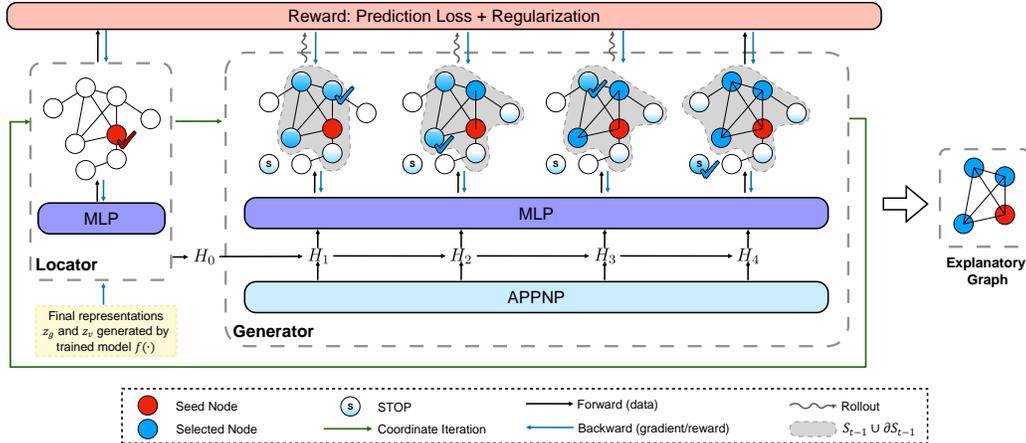


Figure 1: Illustration for explaining GNNs on graph classification. In the inference part, given a graph instance, Locator  $\mathbb{L}$  selects the seed node and Generator  $\mathbb{G}$  generates the explanatory graph step by step until choosing the STOP. In the training part,  $\mathbb{L}$  takes the representation  $z_g$  and  $z_v$  as input, and coordinates a better MLP for the current generator. Given a seed node,  $\mathbb{G}$  learns the parameters in MLP and APPNP to maximize the reward.  $\mathbb{L}$  and  $\mathbb{G}$  train coordinately to optimize the objective.

89 optimization problems. For example, S2V-DQN [6] uses deep Q-learning with graph embedding  
 90 to learn effective algorithms for the Minimum Vertex Cover, the Maximum Cut and the Traveling  
 91 Salesman problems. A graph pointer network is proposed in [18] to solve the TSP efficiently.  
 92 Further, Seal [36] learns heuristics to detect communities in the graph with policy gradient. Note that  
 93 explaining GNNs is also a combinatorial optimization problem. Thus, in this paper, we propose a  
 94 RL-based framework with three dedicated steps to generate explanations.

95 **Post-hoc Analysis in Graph Neural Networks.** By extending existing image/text explanation  
 96 techniques to the graph, some gradient-based methods [22, 1] are proposed to study the importance of  
 97 nodes and edges in the graph. However, their performances have been proved to be sub-optimal [31]  
 98 because they cannot incorporate the special properties of graphs.

99 GNNExplainer [31] is the first specific method proposed to explain trained GNNs. It defines the  
 100 problem as an optimization task, which maximizes the mutual information between the predicted  
 101 labels and the distribution of possible subgraphs under some constraints. Following the problem  
 102 setting, PGExplainer [16] leverages the representations generated by the trained GNN and adopts  
 103 a deep neural network to learn the crucial nodes/edges. These methods both utilize the continuous  
 104 relaxation on edges, and add size and entropy constraints to make the explanation small and sparse.  
 105 Specifically, they optimize a soft mask matrix for edges, and select crucial nodes/edges by the  
 106 threshold. However, they compute the importance of each edge independently, which may lead  
 107 to a disconnected explanatory graph with information redundancies. Our model sequentially adds  
 108 important nodes from the neighbors of the current generated graph, which considers the information  
 109 already involved in the current graph and ensures the connectivity.

110 SubgraphX [34] uses Monte Carlo tree search and Shapley value as a score function to find the  
 111 best connected subgraphs as explanations for GNNs. Causal Screening [28] is another search-based  
 112 method, but it uses greedy search and causality measure to generate the explanations. Different from  
 113 the search-based methods where heuristics are usually hand-crafted, our method uses RL to learn  
 114 heuristics from data, which can be widely applicable. Besides, our learning-based method could train  
 115 by a small set of instances, and infer the explanations of many other similar unseen instances much  
 116 faster than the search-based methods.

117 Different from the instance-level explanation, there also exists the model-level explanation to in-  
 118 vestigate general patterns for predictions. For example, XGNN [33] utilizes the graph generator to  
 119 interpret GNNs at the model-level. In particular, the instance-level explainer interprets the prediction  
 120 for a certain given instance while the model-level explainer is input-independent and less precise.

### 121 3 Preliminary

122 In this section, we first introduce the notations used and then give the formal problem definition.

123 **Graph.** Let  $G = (\mathcal{V}, \mathcal{E})$  denote the graph with node set  $\mathcal{V} = \{v_1, v_2 \cdots v_N\}$  and edge set  $\mathcal{E} \in \mathcal{V} \times \mathcal{V}$ .  
 124 Nodes in  $\mathcal{V}$  could be associated with  $d$ -dimensional node features  $\mathcal{X} \in \mathbb{R}^{N \times d}$ . The graph  $G$   
 125 is described by the adjacency matrix  $A$  such that each entry  $A_{ij} = 1$  if  $e_{ij} \in \mathcal{E}$ ; 0, otherwise.  
 126  $\tilde{A} = A + I_N$  denotes the adjacency matrix with added self-loops. A symmetrically normalized  
 127 adjacency matrix with self-loops  $\hat{A}$  could be computed by  $\tilde{D}^{-1/2} \tilde{A} \tilde{D}^{-1/2}$ , where  $\tilde{D}$  is the diagonal  
 128 degree matrix of  $\tilde{A}$ .

129 **Message Passing.** Given an input graph  $G$  and node features  $\mathcal{X}$ , a GNN model  $f(G, \mathcal{X})$  learns  
 130 node representations. To fuse the information of both node features and graph topology in node  
 131 representation vectors, GNN models utilize the message passing scheme to aggregate information  
 132 from node neighbors. At each layer  $l$ , a node  $v$  aggregates all the messages in the  $(l-1)$ -th layer from  
 133 all its neighbors to generate its embedding:  $h_v^l = \text{update}(\sum_{w \in N(v)} \text{message}(h_v^{l-1}, h_w^{l-1}), h_v^{l-1})$ ,  
 134 where  $N(v)$  is the neighbor set of node  $v$ . With  $L$  layers, each node  $v$  generates its embedding vector  
 135  $z_v = h_v^L$  from the  $L$ -hop neighborhood. In the node classification task, we train a classifier with  $z_v$   
 136 as input to predict the label for node  $v$ . In the graph classification task, we first use an aggregation  
 137 function  $\text{readout}(\{z_v\})$  to generate the graph representation, which is further fed into a classifier to  
 138 predict the label.

139 **Problem Definition.** Given an input graph  $G = (\mathcal{V}, \mathcal{E})$  and a trained GNN model  $f(\cdot)$ , GNN  
 140 explainers aim to generate an explanatory subgraph  $S = (\mathcal{V}_S, \mathcal{E}_S)$ , where  $\mathcal{V}_S \in \mathcal{V}$  and  $\mathcal{E}_S \in \mathcal{E}$ . The  
 141 goal is to maximize the mutual information between the original label prediction  $Y = f(G)$  and the  
 142 label prediction distribution based on the generated explanatory subgraph. Formally, the objective can  
 143 be given as  $\max_S \text{MI}(Y, S) = H(Y) - H(Y|S)$ , where  $\text{MI}(\cdot)$  is the mutual information function  
 144 and  $H(\cdot)$  is the entropy function. Since  $H(Y)$  is fixed in the explanation stage, the objective can be  
 145 rewritten as  $\min_S H(Y|S)$ .

### 146 4 Methodology

147 The objective  $\min_S H(Y|S)$  is intractable since there are exponential candidates for  $S$ . It can be  
 148 considered as a combinatorial optimization problem, where we need to choose a subset of nodes in  
 149 the graph to optimize the objective. We explore how RL can be used to iteratively understand the  
 150 representations produced by GNNs, and generate the explanatory subgraph optimizing the objective.  
 151 Three proposed components will be described in the following.

#### 152 4.1 Iterative Graph Generation (Graph Generator)

153 Given a starting point  $v_0$ , the graph generator is used to generate a connected subgraph  $S =$   
 154  $\{v_0, v_1 \cdots v_T\}$ , where we select one node in a step and  $T$  is the total number of steps. Specifically,  
 155 at the  $t$ -th step, we have the current partial solution  $S_t = \{v_0, v_1 \cdots v_{t-1}\}$ . We next select a  
 156 new node  $v_t$  from the boundary  $\partial S_{t-1}$  and expand the solution  $S_t = S_{t-1} \cup \{v_t\}$ . The state is  
 157 defined as the combined representation of both  $v_0$  and  $S_{t-1}$ . The action space is the boundary  
 158  $\partial S_{t-1} = \cup_{v \in S_{t-1}} N(v) \setminus S_{t-1}$ . We further associate the solution with a reward value.

159 **State.** At the  $t$ -th step, we first augment each node feature vector by adding the information of the  
 160 starting point and the current partial subgraph. For each node  $v$  in  $S_{t-1}$  or  $\partial S_{t-1}$ , we concatenate  
 161 two values with its original feature vector  $x_v$ :

$$x'_v = [x_v, \mathbb{1}_{\{v \in \{v_0\}\}}, \mathbb{1}_{\{v \in S_{t-1}\}}], \quad X'_t = [x'_v]_{\forall v \in S_{t-1} \cup \partial S_{t-1}}, \quad (1)$$

162 where  $\mathbb{1}$  is the indicator function and  $\mathbb{1}_{\{v \in S\}} = 1$  if  $v \in S$  otherwise 0. We concatenate the  
 163 augmented node features and obtain the initial state representation  $X'_t$ .

164 Each node could further combine information from its current neighborhood. To achieve this, we  
 165 utilize some existing GNN methods, e.g., APPNP [13], which separate the non-linear transformation

166 and information propagation. These methods have shown to be highly efficient and also effective.  
 167 Specifically, we have the following update equation:

$$H_t^{(0)} = \Theta_1 X_t', \quad H_t^{(l+1)} = (1 - \alpha) \hat{A} H_t^{(l)} + \alpha H_t^{(0)}, \quad (2)$$

168 where  $\Theta_1$  is the trainable weight matrix,  $\hat{A}$  is the symmetrically adjacency matrix, and  $\alpha$  is a hyper-  
 169 parameter used to control weight. After  $L$ -layer updates, we obtain the node representations  $H_t^L$ . we  
 170 feed them into a MLP to improve the representation ability:

$$\bar{H}_t(v) = \text{MLP}(H_t^L(v); \Theta_2), v \in S_{t-1} \cup \partial S_{t-1}, \quad (3)$$

171 where  $\Theta_2$  is the learnable parameters in the MLP. At the  $t$ -th step, we only consider the nodes in  $S_{t-1}$   
 172 and  $\partial S_{t-1}$  because other nodes do not influence the state and action space.

173 **Action.** Since the connectivity of the generated subgraph is required, we take  $\partial S_{t-1}$  as the action  
 174 space at the  $t$ -th step. We utilize a softmax function to calculate the probability of taking an action  
 175  $v \in \partial S_{t-1}$ :

$$a_t(v) = \frac{\exp(\theta_3^T \bar{H}_t(v))}{\sum_{u \in \partial S_{t-1}} \exp(\theta_3^T \bar{H}_t(u))}, v \in \partial S_{t-1}, \quad (4)$$

176 where  $\theta_3$  is the trainable parameter vector.

177 **Objective.** Following [16], we use the cross-entropy function to replace the conditional entropy  
 178 function  $\min_S H(Y|S)$  with  $N$  given instances. We rewrite the objective as:

$$\text{Prediction Loss} = -\frac{1}{N} \sum_{n=1}^N \sum_{c=1}^C P(Y=c) \log P(f(S_n)=c), \quad (5)$$

179 where  $S_n$  is the explanatory subgraph for the  $n$ -th instance,  $C$  is the number of possible predicted  
 180 labels,  $P(Y=c)$  is the probability that the original output of the trained GNN  $f$  is  $c$ , and  $P(f(S_n)=$   
 181  $c)$  is the probability that the label prediction of  $f$  on the subgraph  $S_n$  is  $c$ .

182 We further introduce some regularization terms to restrict the characteristics of the explanatory  
 183 subgraph. To obtain a compact and succinct explanatory subgraph  $S$ , we define a size loss and a  
 184 radius penalty, respectively. The size loss is used to limit the number of nodes in  $S$  while the radius  
 185 penalty can compute the longest length of the shortest path from the seed node to other nodes in  $S$ . We  
 186 also introduce a similarity loss that measures the similarity between the original node representation  
 187  $z_{v_0}$  and the new representation generated on  $S$ . Formally, these loss function are defined as:

$$\text{Size Loss} = \|S\|_1, \quad \text{Radius Penalty} = \max_{u \in S} \text{Distance}(v_0, u), \quad \text{Similarity Loss} = \|\bar{H}_T(v_0) - z_{v_0}\|_2. \quad (6)$$

188 The final objective is to minimize

$$\mathcal{L}(S) = \text{Prediction Loss} + \lambda_1 \cdot \text{Size Loss} + \lambda_2 \cdot \text{Radius Penalty} + \lambda_3 \cdot \text{Similarity Loss}, \quad (7)$$

189 where  $\lambda_*$  are hyper-parameters to control the term importance.

190 **Reward.** We take the objective function loss  $\mathcal{L}$  as the (negative) reward. However, the loss cannot  
 191 be separated into each generation step. If we simply compute  $-\mathcal{L}(S_t)$  at the  $t$ -th step and regard it as  
 192 the reward  $r_t$  for the state-action pair  $(s_t, a_t)$ , it could lead to the sub-optimal results. Therefore, we  
 193 do not compute intermediate rewards when adding a new node to the subgraph. We only return the  
 194 reward  $-\mathcal{L}(S)$  when we complete the generation process of  $S$ .

195 **Optimization with policy gradient.** We learn the graph generator  $\mathbb{G} = \{\Theta_1, \Theta_2, \theta_3\}$  via policy  
 196 gradient. The policy  $\pi_\theta$  is learned to maximize  $\mathbb{E}_{S|v_0 \sim \mathbb{G}}[-\mathcal{L}(S)]$ , whose policy gradient is

$$\nabla \mathbb{E}_{S|v_0 \sim \mathbb{G}}[-\mathcal{L}(S)] = \mathbb{E}_{v_1, \dots, v_T | v_0 \sim \mathbb{G}} \left[ \sum_{t=1}^T \nabla \log \pi_\theta(v_t | S_{t-1}) \cdot Q(S_{t-1}, v_t) \right]. \quad (8)$$

197 Here,  $S_0$  is the given seed  $v_0$  and  $Q(S_{t-1}, v_t) = \mathbb{E}_{v_{t+1}, \dots, v_T | S_{t-1} \cup v_t \sim \mathbb{G}}[-\mathcal{L}(S)]$  is the state-action value  
 198 function. Specifically, we use the Monte-Carlo estimation to approximate  $Q$  values:

$$Q(S_{t-1}, v_t) = \begin{cases} \frac{1}{M} \sum_{i=1}^M -\mathcal{L}(S^{(i)}) & \text{when } t < T \\ -\mathcal{L}(S_{t-1} \cup v_t) & \text{when } t = T \end{cases} \quad (9)$$

199 where  $S^{(i)}$  ( $i = 1, \dots, M$ ) are rollouts (i.e., complete explanatory graphs) sampled from the policy  
 200 given the partial solution  $S_{t-1}$  and  $v_t$ . In this way, we can solve the sparse reward problem and also  
 201 distribute reward signals at all steps.

202 **4.2 Stopping Criteria Learning**

203 We further learn the stopping criteria to judge the goodness of the current generated graph. We add a  
 204 special STOP action into the action space to learn node selection and stopping criteria simultaneously.  
 205 Since we have already obtained the node representation  $\bar{H}_t$  in the current state, the STOP action  
 206 could aggregate the representations with self-attention mechanism:

$$\gamma_t(v) = \frac{\exp(\theta_4^T \bar{H}_t(v))}{\sum_{u \in \{S_{t-1} \cup \partial S_{t-1}\}} \exp(\theta_4^T \bar{H}_t(u))}, \quad \bar{H}_t(\text{STOP}) = \sum_{v \in \{S_{t-1} \cup \partial S_{t-1}\}} \gamma_t(v) \bar{H}_t(v). \quad (10)$$

207 where the parameter  $\theta_4$  helps to learn the attention  $\gamma_t(v)$  for each node  $v$  in the current state.

208 After having  $\bar{H}_t(\text{STOP})$ , we could put it into Eqn. 4, and compute  $a_t(\text{STOP})$  and  $a_t(v)_{\forall v \in \partial S_{t-1}}$   
 209 together. In practice, we also set a maximum number of generation steps to avoid generating very  
 210 large subgraphs.

211 **4.3 Starting Point Selection (Seed Locator)**

212 For node classification tasks, the starting point is the node instance whose predicted label needs to  
 213 be interpreted. However, the starting point is difficult to select for a graph instance. To solve the  
 214 problem, we need to construct a seed locator  $\mathbb{L}$  to first identify the most influential node in the graph  
 215 and then generate the explanatory subgraph from that node.

216 Given  $N$  graph instances  $g_n$ , the objective in Eqn. 5 could be rewritten with the locator  $\mathbb{L}$ :

$$\min_{\mathbb{G}, \mathbb{L}} -\frac{1}{N} \sum_{n=1}^N \sum_{c=1}^C P(f(g_n) = c) \log P(f(\mathbb{G}(\mathbb{L}(g_n))) = c), \quad (11)$$

217 where both the generator  $\mathbb{G}$  and the locator  $\mathbb{L}$  are to be learned. The regularization terms in Eqn. 6 can  
 218 be further added to the objective for more constraints. We train  $\mathbb{G}$  and  $\mathbb{L}$  coordinately, i.e., we fix the  
 219 parameters in one module and train the other module to optimize the objective in Eqn. 11 iteratively.

220 When we fix  $\mathbb{L}$ , the way to train  $\mathbb{G}$  is the same as described in Sec. 4.1. Here we introduce how to  
 221 construct  $\mathbb{L}$  when  $\mathbb{G}$  is fixed. Based on  $\mathbb{G}$ , we can generate a subgraph  $S$  for each node in the graph  
 222 instance and compute the corresponding reward  $-\mathcal{L}(S)$ . A straightforward way is to enumerate over  
 223 all the nodes and select the node with the highest reward. However, such a brute-force method is  
 224 computationally infeasible when we have many graph instances. Therefore, we adopt a learning-based  
 225 method. Specifically, we use a three-layer MLP to model the influence of a node  $v_{i,n}$  on the label of  
 226 the graph instance  $g_n$ :

$$\omega_{i,n} = \text{MLP}([z_{g_n}, z_{v_{i,n}}]), \quad (12)$$

227 where  $z_{g_n}$  and  $z_{v_{i,n}}$  are the final representations of the graph instance  $g_n$  and the node  $v_{i,n}$  generated  
 228 by the trained model  $f(\cdot)$ , respectively. Because the goal of  $\mathbb{L}$  is to return the node  $v_{i,n}$  with highest  
 229  $\omega_{i,n}$  for graph  $g_n$ , we utilize the Kullback-Leibler divergence loss,  $\text{KLDivLoss}(\omega_{i,n}, -\mathcal{L}(\mathbb{G}(v_{i,n})))$ ,  
 230 which aims to make the distribution between estimated values  $\omega_{i,n}$  and the actual reward of explana-  
 231 tory subgraph produced by the current generator close. The softmax layers are used to transform  $\omega_{i,n}$   
 232 and  $-\mathcal{L}(\mathbb{G}(v_{i,n}))$  into two distributions over the nodes  $v_{i,n}$  in graph  $g_n$ . We sample graph instances  
 233 to train the MLP, and let  $\mathbb{L}$  learn what kind of seed nodes has the highest reward to minimize Eqn.11.

234 **4.4 Pre-training**

235 In this section, we show the pre-training strategies with Maximum Log-Likelihood Estimation (MLE)  
 236 that can be used to initialize the generator  $\mathbb{G}$  and the locator  $\mathbb{L}$ , respectively.

237 **Pretrain  $\mathbb{G}$ .** The generator  $\mathbb{G}$  produces an unordered set  $S$  but in an ordered sequence. We maximize  
 238 over all possible generated orderings for an explanatory graph [26]:

$$\max_{\tau} \sum_{i=1}^T \log \mathbb{G}(v_{\tau(i)} | \{v_0, v_{\tau(1)}, \dots, v_{\tau(i-1)}\}), \quad (13)$$

239 where  $\tau$  is any *valid* permutation given  $S$ . Here, the validity means that each  $v_{\tau(i)}$  should be in the  
 240 boundary of  $\{v_0, v_{\tau(1)}, \dots, v_{\tau(i-1)}\}$ .

241 Considering there are at most  $T!$  orderings in the worst case, we use a bootstrapped way with  
 242 set2set [26] to approximate Eqn. 13. Specifically, we optimize Eqn. 13 by maximizing the following  
 243 set-wise log-likelihood instead:

$$\sum_{i=1}^T \log \mathbb{G}(v_{\tau^*(i)} | \{v_0, v_{\tau^*(1)}, \dots, v_{\tau^*(i-1)}\}), \quad (14)$$

244 where  $\tau^*(i) = \arg \max \mathbb{G}(\cdot | \{v_0, v_{\tau^*(1)}, \dots, v_{\tau^*(i-1)}\})$ . Here we utilize the 3-hop neighborhood of  
 245 seed nodes as the initial explanations  $S$  (i.e., pre-training samples). After constructing the samples,  
 246 we train  $\mathbb{G}$  to maximize the set-wise log-likelihood in Eqn. 14. Vinyals et al. [26] also pointed out if  
 247 we naively optimize it, the model would pick a random ordering and get stuck on it. Thus, a list-wise  
 248 log-likelihood is also necessary to explore the space of ordering. For each pretrain sample  $S$ , we  
 249 choose a valid permutation  $\tau'$  with random lengths in advance, and optimize the list-wise MLE:  
 250  $\sum_{i=1}^T \log \mathbb{G}(v_{\tau'(i)} | \{v_0, v_{\tau'(1)}, \dots, v_{\tau'(i-1)}\})$ .

251 **Pretrain  $\mathbb{L}$ .** We pretrain the locator  $\mathbb{L}$  without the generator  $\mathbb{G}$ . Similar as in  $\mathbb{G}$ , we utilize the  
 252 3-hop neighborhood of a node as the initial explanatory subgraph  $S$ . We randomly sample some  
 253 nodes in the graph instances and compute the rewards of their 3-hop neighborhoods. These samples  
 254 are used to pretrain the parameters in  $\mathbb{L}$ .

## 255 5 Experiments

256 In this section, we first introduce our experimental setup. Then we compare RG-Explainer with  
 257 two state-of-the-art baselines GNNExplainer [31] and PGExplainer [16] in both qualitative and  
 258 quantitative evaluations. Further, we evaluate the performance of our method in the inductive setting.  
 259 Due to the space limitation, we move the pseudocode, implementation details and ablation study to  
 260 the supplementary materials. We also attach our codes in the supplementary materials.

### 261 5.1 Setup

262 For fairness, we follow the experimental setup in [16, 11], i.e., the same datasets, trained GNN  
 263 model and evaluation metrics. Besides, we also utilize the same fine-tuned parameters in [11] for our  
 264 competitors, GNNExplainer and PGExplainer.

265 **Datasets.** We use six datasets, in which four synthetic datasets (BA-shapes, BA-Community, Tree-  
 266 Cycles and Tree-Grid) are used for the node classification task and two datasets (BA-2motifs and  
 267 Mutagenicity) are used for the graph classification task. These datasets are composed of *motifs* and  
 268 *bases*. The motif is a small but important substructure in a graph, which has been shown to play a  
 269 crucial role in predicting the label of node/graph instances [5, 14, 15]. The base is the remaining  
 270 part of a graph which is randomly generated. Motifs are taken as the ground-truth and the goal of  
 271 explainers is to find them. Details of these datasets are described as follows.

272 (a) The BA-shapes dataset consists of one Barabasi-Albert(BA) graph [2] as the base and 80 house-  
 273 structure motifs. Each motif is randomly attached to a node in BA graph and extra edges are added as  
 274 noises. (b) The BA-community dataset is comprised of two BA-shapes with different node features  
 275 generated by Gaussian distributions. The extra edges are also added to connect two BA-shapes. (c)  
 276 The Tree-cycles dataset includes a multi-level binary tree as the base and 80 six-node cycle motifs.  
 277 The cycle motifs are randomly attached to the tree. (d) The Tree-grid dataset is similar to Tree-cycles,  
 278 which uses the  $3 \times 3$  grid motifs instead. (e) The BA-motifs dataset has 1000 graphs where half of  
 279 them are a BA graph attached with a house-structure motif, while the rest are a BA graph attached  
 280 with a five-node cycle motif. (f) The Mutagenicity dataset is a real dataset, which includes 4337  
 281 molecule graphs. They can be classified as mutagenic or nonmutagenic depending on whether having  
 282  $NH_2$  or  $NO_2$  motifs.

283 **Model.** We use the trained GNN model in [11], whose architecture is given in [16, 31]. Specifi-  
 284 cally, the model that consists of three consecutive Graph Convolution layers connected with a fully  
 285 connected layer is used for node classification. For graph classification, the model includes three  
 286 consecutive Graph Convolution layers fed into two max and mean pooling layers, respectively. The  
 287 two pooling layer output embeddings are then concatenated to generate the input for a fully connected  
 288 layer for graph classification.

Table 1: Visualization (Qualitative Evaluation)

	Node Classification				Graph Classification	
	BA-Shapes	BA-Community	Tree-Cycles	Tree-Grid	BA-2motifs	MUTAG
Explanations by GNN-Explainer						
Explanations by PG-Explainer						
Explanations by RG-Explainer (ours)						
Ground-Truth Motif						

289 **Metrics.** The motifs in each dataset are the ground-truth explanations. The edges in the motif  
 290 are positive and other edges are negative. GNNExplainer and PGExplainer return a mask matrix  
 291 to represent the importance of each edge in the instance. Our method generates a subgraph. Based  
 292 on the generation order of edges, we could also assign different weights to edges. Therefore, the  
 293 explanation problem can be formalized as a binary classification task, where edges in the ground-truth  
 294 motif are taken as prediction labels and the weights of edges are viewed as prediction scores. With  
 295 the explanatory subgraph provided by explainers, the **AUC score** can be computed to measure the  
 296 accuracy for quantitative evaluation.

## 297 5.2 Qualitative evaluation

298 Table 1 visualizes some examples of explanatory graphs on all the datasets. For node classification,  
 299 we amplify the center node instance and generate the subgraph from it. For graph classification, the  
 300 graph represents the whole graph instance. We use different colors to denote node labels.

301 Intuitively, a superior explainer should include more edges in the ground-truth motif and less irrelevant  
 302 nodes and edges in the explanatory subgraph. To show whether explainers assign larger weights to  
 303 edges in the motif, we use the bold black edges to represent the top- $k$  edges, where  $k$  is the number  
 304 of edges inside the ground-truth motif.

305 For easy cases, all the methods could find the ground-truth motif. Thus, we choose some difficult  
 306 instances to show advantages of our method. From the table, we see that all the three methods  
 307 can generate subgraphs that contain the ground-truth “house” motif on both BA-Shapes and BA-  
 308 Community datasets. However, both GNNExplainer and PGExplainer include many irrelevant nodes  
 309 and edges, while our method adopts the stopping criteria to generate more concise subgraphs. On  
 310 both Tree-Cycles and Tree-Grid datasets, since we select the connection node between the base and  
 311 the motif as the node instance, it is hard to identify the ground-truth motif (cycle or grid) exactly.  
 312 This is because the explanatory subgraphs could easily include the base. Compared with competitors,  
 313 our method can generate the explanatory subgraphs that contain the complete ground-truth motifs,  
 314 i.e., edges in the motifs are all marked as black.

315 For graph classification, explainers should identify the ground-truth motif that decides the label of  
 316 graph instances. For the instance in BA-2motifs, the ground-truth “house” motif is located at the right  
 317 bottom. Both GNNExplainer and PGExplainer involve accurate and wrong edges in the explanations,  
 318 while our method only adds edges in the ground-truth “house” motif to the generated subgraph. For  
 319 the MUTAG dataset, both PGExplainer and our method identify the correct motif. In particular, our  
 320 method correctly locates the  $O$  atom and obtains the ground-truth motif  $NO_2$ .

Table 2: Explanation AUC (Quantitative Evaluation).

	Node Classification				Graph Classification	
	BA-Shapes	BA-Community	Tree-Cycles	Tree-Grid	BA-2motifs	MUTAG
GNNExplainer	0.742 ± 0.006	0.708 ± 0.004	0.540 ± 0.017	0.714 ± 0.002	0.499 ± 0.004	0.606 ± 0.003
PGExplainer	<b>0.999 ± 0.000</b>	0.825 ± 0.040	0.760 ± 0.014	0.679 ± 0.008	0.133 ± 0.046	0.847 ± 0.081
RG-Explainer (ours)	0.985 ± 0.013	<b>0.919 ± 0.017</b>	<b>0.787 ± 0.099</b>	<b>0.927 ± 0.031</b>	<b>0.657 ± 0.107</b>	<b>0.873 ± 0.028</b>
Improve	-1.5%	11.4%	3.6%	29.8%	31.7%	2.8%

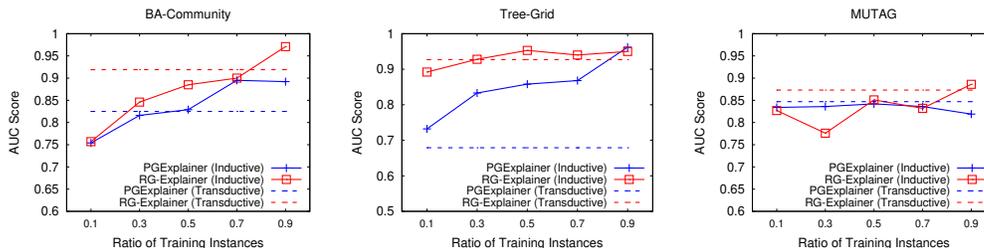


Figure 2: Comparison between RG-Explainer and PGExplainer in the inductive setting.

### 321 5.3 Quantitative evaluation

322 We next show the quantitative results in Table 2. For each method, we compute the average AUC  
 323 scores and standard deviations over 10 runs. Note that the AUC scores reported here are a bit different  
 324 from that in the original papers due to the unstable convergences [11].

325 From the table, we see that our method RG-Explainer achieves the best results on 5 out of 6 datasets.  
 326 For example, the AUC score of RG-Explainer on Tree-Grid is 0.927 while that of the runner-up  
 327 is only 0.714, leading to an improvement of 29.8%. On the BA-Shapes dataset, RG-Explainer  
 328 achieves comparable results with the winner’s and significantly outperforms GNNExplainer. These  
 329 results show the advantage of applying reinforcement learning techniques in constructing explanatory  
 330 subgraphs. Further, since the constructed subgraphs are connected, they could better characterize  
 331 motifs in the graph. Note that in the graph classification task, our method uses a locator to first select  
 332 the seed nodes. We also test the performance of the locator and find that the locator selects  $\sim 66\%$   
 333 and  $\sim 84\%$  accurate seed nodes (i.e., nodes in the ground-truth motif) for BA-2Motifs and MUTAG,  
 334 respectively. This further explains the good performance of RG-Explainer for graph classification.

### 335 5.4 Inductive setting

336 We further test the performance of RG-Explainer in the inductive setting. We compare it with  
 337 PGExplainer, which are both learning-based methods. Specifically, we vary the training set sizes  
 338 from  $\{10\%, 30\%, 50\%, 70\%, 90\%\}$  and take the remaining instances for testing. For each dataset,  
 339 we run the experiments 10 times and compute the average AUC scores.

340 Due to the limited space, Fig 2 only shows the results on BA-Community, Tree-Grid and MUTAG  
 341 datasets. For other datasets, see the supplementary material. The figure also includes the results of  
 342 both methods in the transductive setting (i.e., use all the instances) for reference. From the figure,  
 343 RG-Explainer generally outperforms PGExplainer as the training set size increases. For example,  
 344 with only 10% training instances in the Tree-Grid dataset, RG-Explainer significantly outperforms  
 345 PGExplainer by a large margin. This shows that RG-Explainer generalizes better than PGExplainer.

## 346 6 Conclusion

347 We present RG-Explainer to generate the instance-level explanations for GNNs in this paper. The  
 348 RL-based generator is proposed to ensure the message passing nature of GNNs. Besides, we design  
 349 the seed locator and stopping criteria to find the most influential node in a graph instance and check  
 350 whether the generated explanatory graph is good enough, respectively. Though our method increases  
 351 the transparency of GNN predictions, it may put GNN models at a high risk of being attacked. How  
 352 to utilize the GNN explanations to make GNN models more robust is a future research direction.

## 353 References

- 354 [1] F. Baldassarre and H. Azizpour. Explainability techniques for graph convolutional networks.  
355 *arXiv preprint arXiv:1905.13686*, 2019.
- 356 [2] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *science*,  
357 286(5439):509–512, 1999.
- 358 [3] D. Barrett, F. Hill, A. Santoro, A. Morcos, and T. Lillicrap. Measuring abstract reasoning in  
359 neural networks. In *ICML*, pages 511–520. PMLR, 2018.
- 360 [4] P. W. Battaglia, R. Pascanu, M. Lai, D. Rezende, and K. Kavukcuoglu. Interaction networks for  
361 learning about objects, relations and physics. *arXiv preprint arXiv:1612.00222*, 2016.
- 362 [5] A. R. Benson, D. F. Gleich, and J. Leskovec. Higher-order organization of complex networks.  
363 *Science*, 353(6295):163–166, 2016.
- 364 [6] H. Dai, E. B. Khalil, Y. Zhang, B. Dilkina, and L. Song. Learning combinatorial optimization  
365 algorithms over graphs. In *NeurIPS*, pages 6348–6358, 2017.
- 366 [7] N. De Cao and T. Kipf. Molgan: An implicit generative model for small molecular graphs.  
367 *arXiv preprint arXiv:1805.11973*, 2018.
- 368 [8] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural message passing for  
369 quantum chemistry. In *ICML*, pages 1263–1272, 2017.
- 370 [9] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and  
371 Y. Bengio. Generative adversarial nets. In *NeurIPS*, pages 2672–2680, 2014.
- 372 [10] W. Hamilton, Z. Ying, and J. Leskovec. Inductive representation learning on large graphs. In  
373 *NeurIPS*, pages 1024–1034, 2017.
- 374 [11] L. Holdijk, M. Boon, S. Henckens, and L. de Jong. [re] parameterized explainer for graph  
375 neural network. In *ML Reproducibility Challenge 2020*, 2021.
- 376 [12] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks.  
377 In *ICLR*, 2017.
- 378 [13] J. Klicpera, A. Bojchevski, and S. Günnemann. Predict then propagate: Graph neural networks  
379 meet personalized pagerank. In *ICLR*, 2019.
- 380 [14] A. S. Konagurthu and A. M. Lesk. On the origin of distribution patterns of motifs in biological  
381 networks. *BMC Systems Biology*, 2(1):1–8, 2008.
- 382 [15] X. Li, R. Cheng, K. C.-C. Chang, C. Shan, C. Ma, and H. Cao. On analyzing graphs with  
383 motif-paths. *Proceedings of the VLDB Endowment*, 14(6):1111–1123, 2021.
- 384 [16] D. Luo, W. Cheng, D. Xu, W. Yu, B. Zong, H. Chen, and X. Zhang. Parameterized explainer  
385 for graph neural network. In *NeurIPS*, 2020.
- 386 [17] Y. Luo, K. Yan, and S. Ji. Graphdf: A discrete flow model for molecular graph generation. In  
387 *ICML*, 2021.
- 388 [18] Q. Ma, S. Ge, D. He, D. Thaker, and I. Drori. Combinatorial optimization by graph pointer  
389 networks and hierarchical reinforcement learning. *arXiv preprint arXiv:1911.04936*, 2019.
- 390 [19] T. Ma, J. Chen, and C. Xiao. Constrained generation of semantically valid graphs via regularizing  
391 variational autoencoders. In *NeurIPS*, 2018.
- 392 [20] K. Madhawa, K. Ishiguro, K. Nakago, and M. Abe. Graphnvp: An invertible flow model for  
393 generating molecular graphs. *arXiv preprint arXiv:1905.11600*, 2019.
- 394 [21] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves,  
395 M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep rein-  
396 forcement learning. *nature*, 518(7540):529–533, 2015.

- 397 [22] P. E. Pope, S. Kolouri, M. Rostami, C. E. Martin, and H. Hoffmann. Explainability methods for  
398 graph convolutional neural networks. In *CVPR*, pages 10772–10781, 2019.
- 399 [23] D. Saxton, E. Grefenstette, F. Hill, and P. Kohli. Analysing mathematical reasoning abilities of  
400 neural models. *arXiv preprint arXiv:1904.01557*, 2019.
- 401 [24] C. Shi, M. Xu, Z. Zhu, W. Zhang, M. Zhang, and J. Tang. Graphaf: a flow-based autoregressive  
402 model for molecular graph generation. In *ICLR*, 2020.
- 403 [25] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio. Graph attention  
404 networks. In *ICLR*, 2018.
- 405 [26] O. Vinyals, S. Bengio, and M. Kudlur. Order matters: Sequence to sequence for sets. *ICLR*,  
406 2015.
- 407 [27] M. N. Vu and M. T. Thai. Pgm-explainer: Probabilistic graphical model explanations for graph  
408 neural networks. In *NeurIPS*, 2020.
- 409 [28] X. Wang, Y. Wu, A. Zhang, X. He, and T. seng Chua. Causal screening to interpret graph neural  
410 networks, 2021.
- 411 [29] F. Wu, A. Souza, T. Zhang, C. Fifty, T. Yu, and K. Weinberger. Simplifying graph convolutional  
412 networks. In *ICML*, pages 6861–6871. PMLR, 2019.
- 413 [30] K. Xu, W. Hu, J. Leskovec, and S. Jegelka. How powerful are graph neural networks? In *ICLR*,  
414 2018.
- 415 [31] R. Ying, D. Bourgeois, J. You, M. Zitnik, and J. Leskovec. Gnnexplainer: Generating explana-  
416 tions for graph neural networks. In *NeurIPS*, pages 9240–9251, 2019.
- 417 [32] J. You, B. Liu, R. Ying, V. Pande, and J. Leskovec. Graph convolutional policy network for  
418 goal-directed molecular graph generation. *arXiv preprint arXiv:1806.02473*, 2018.
- 419 [33] H. Yuan, J. Tang, X. Hu, and S. Ji. Xggn: Towards model-level explanations of graph neural  
420 networks. In *KDD*, pages 430–438, 2020.
- 421 [34] H. Yuan, H. Yu, J. Wang, K. Li, and S. Ji. On explainability of graph neural networks via  
422 subgraph explorations. *ICML*, 2021.
- 423 [35] M. Zhang and Y. Chen. Link prediction based on graph neural networks. In *NeurIPS*, pages  
424 5165–5175, 2018.
- 425 [36] Y. Zhang, Y. Xiong, Y. Ye, T. Liu, W. Wang, Y. Zhu, and P. S. Yu. Seal: Learning heuristics for  
426 community detection with generative adversarial networks. In *KDD*, pages 1103–1113, 2020.

## 427 Checklist

- 428 1. For all authors...
- 429 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s  
430 contributions and scope? [Yes]
- 431 (b) Did you describe the limitations of your work? [Yes] See the conclusion section and  
432 supplementary materials
- 433 (c) Did you discuss any potential negative societal impacts of your work? [Yes] See the  
434 conclusion section and supplementary materials
- 435 (d) Have you read the ethics review guidelines and ensured that your paper conforms to  
436 them? [Yes]
- 437 2. If you are including theoretical results...
- 438 (a) Did you state the full set of assumptions of all theoretical results? [N/A]
- 439 (b) Did you include complete proofs of all theoretical results? [N/A]
- 440 3. If you ran experiments...

- 441 (a) Did you include the code, data, and instructions needed to reproduce the main ex-  
442 perimental results (either in the supplemental material or as a URL)? [Yes] See the  
443 supplementary material.
- 444 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they  
445 were chosen)? [Yes] See the supplementary material.
- 446 (c) Did you report error bars (e.g., with respect to the random seed after running experi-  
447 ments multiple times)? [Yes] See the experimental section. We report AVG and STD  
448 over 10 random seeds.
- 449 (d) Did you include the total amount of compute and the type of resources used (e.g., type  
450 of GPUs, internal cluster, or cloud provider)? [Yes] See the supplementary material.
- 451 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 452 (a) If your work uses existing assets, did you cite the creators? [Yes]
- 453 (b) Did you mention the license of the assets? [N/A] The repos we used do not mention  
454 the license.
- 455 (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
- 456 (d) Did you discuss whether and how consent was obtained from people whose data you're  
457 using/curating? [No]
- 458 (e) Did you discuss whether the data you are using/curating contains personally identifiable  
459 information or offensive content? [N/A]
- 460 5. If you used crowdsourcing or conducted research with human subjects...
- 461 (a) Did you include the full text of instructions given to participants and screenshots, if  
462 applicable? [N/A]
- 463 (b) Did you describe any potential participant risks, with links to Institutional Review  
464 Board (IRB) approvals, if applicable? [N/A]
- 465 (c) Did you include the estimated hourly wage paid to participants and the total amount  
466 spent on participant compensation? [N/A]