Soft-Label Integration for Robust Toxicity Classification

Zelei Cheng
Northwestern University
Evanston, USA
zelei.cheng@northwestern.edu

Xian Wu
Northwestern University
Evanston, USA
xianwu2024@u.northwestern.edu

Jiahao Yu
Northwestern University
Evanston, USA
jiahao.yu@northwestern.edu

Shuo Han
Northwestern University
Evanston, USA
shuo.han.1@u.northwestern.edu

Xin-Qiang Cai
The University of Tokyo
Tokyo, Japan
xinqiang.cai@riken.jp

Xinyu Xing
Northwestern University
Evanston, USA
xinyu.xing@northwestern.edu

Abstract

Disclaimer. This paper contains uncensored toxic content that might be offensive.

Toxicity classification in textual content remains a significant problem. Data with labels from a single annotator fall short of capturing the diversity of human perspectives. Therefore, there is a growing need to incorporate crowdsourced annotations for training an effective toxicity classifier. Additionally, the standard approach to training a classifier using empirical risk minimization (ERM) may fail to address the potential shifts between the training set and testing set due to exploiting spurious correlations. This work introduces a novel bi-level optimization framework that integrates crowdsourced annotations with the soft-labeling technique and optimizes the soft-label weights by Group Distributionally Robust Optimization (GroupDRO) to enhance the robustness against out-of-distribution (OOD) risk. We theoretically prove the convergence of our bi-level optimization algorithm. Experimental results demonstrate that our approach outperforms existing baseline methods in terms of both average and worst-group accuracy, confirming its effectiveness in leveraging crowdsourced annotations to achieve more effective and robust toxicity classification.

1 Introduction

Large language models (LLMs) are rapidly being adopted in applications such as conversations [1, 2], AI-assisted programming [3], and education [4]. However, despite impressive capabilities, the interaction between humans and LLMs can generate harmful, biased, or factually incorrect content [5, 6]. For example, users may ask LLMs to generate toxic content, such as hate speech, misinformation, or violent threats, which can have severe consequences for individuals and communities. Recent studies on jailbreaking LLMs also show that adversarial prompts can elicit toxic responses from models [7, 6, 8, 9]. Therefore, there is a pressing need to develop a robust toxicity classification model that can effectively identify and mitigate harmful content generated by LLMs.

Figure 1: **An example of a toxic response with the spurious feature "I must remind you that".** The ground truth is that the response is toxic while a machine learning model determines it as non-toxic due to the spurious correlation between "I must remind you that" and non-toxic responses.

Traditional toxicity classification methods [10–13], typically reliant on labels from a single annotator per instance, fall short of capturing the diversity of human perspectives [14]. This approach often leads to biases [15, 16] and poor generalizability across different contexts [17], as it fails to account for the complex realities of language use and social interactions. Thus, there is a growing need to incorporate crowdsourced annotations that reflect a broader array of cultural and linguistic nuances. Additionally, Arjovsky et al. [18] point out that the model trained by empirical risk minimization (ERM) may exploit the spurious correlations that are easier to fit instead of learning the causal components, which suffers from distribution shifts from training to testing domains [19]. When *spurious correlations* are present, the performance of certain groups of examples can drop significantly. For example, the toxicity classifier might learn to associate certain phrases or contexts (*e.g.,* "I must remind you that" in Figure 1) with non-toxic behavior, despite the overall response being harmful.

To overcome the above challenges, we propose a bi-level optimization framework that incorporates crowdsourced annotations through soft-labeling techniques to enhance the robustness and reliability of toxicity classification systems. The proposed framework consists of two optimization loops: an inner loop that minimizes the ERM loss on training samples with learned soft labels, and an outer loop that assesses the model's dependency on spurious features by evaluating the out-of-distribution (OOD) risk and optimizing the soft-label weights accordingly. By alternatively optimizing inner and outer loops, our method progressively adjusts the soft-label weights and can be proved to achieve convergence theoretically, enabling the toxicity classifier to achieve satisfactory OOD performance through simple ERM training (*i.e.,* inner loop optimization).

Empirically, we evaluate our method on the toxic question classification and response classification datasets provided by a third-party security company and the public HateXplain dataset [20]. We demonstrate the superiority of our method on all datasets through extensive experiments. Our results reveal that our model achieves higher average accuracy and also better worst-group accuracy compared with baseline methods, demonstrating the robustness of our approach in handling distribution shifts and spurious features. Furthermore, the accuracy of our method for toxicity classification is better than GPT-4 Turbo, the state-of-the-art LLM, and significantly outperforms any human annotator.

By integrating multiple annotations and adopting a robust optimization approach, our study not only advances the technological frontiers of toxicity classification but also contributes to the broader discourse on ethical AI practices, promoting more nuanced and equitable online interactions.

2 Related Work

Bi-level Optimization. Bi-level optimization [21] has attracted significant attention due to its ability to handle hierarchical decision-making tasks including meta learning [22–26], neural architecture search [27–29], sample re-weighting [30, 31, 25], label denoising [32], etc. For example, in meta-learning [26], bi-level optimization provides a way to learn the initial parameters of a model which leads to fast adaptation and good generalization for various learning tasks. In this work, we formulate the toxicity classification from multiple annotations as a bi-level optimization problem where we alternate between minimizing the empirical risk minimization (ERM) loss on training samples with learned soft labels and optimizing the soft-label weights against the out-of-distribution (OOD) risk.

Learning from Partial Labels. Training a classifier from partial labels implicitly requires determining the ground truth from multiple annotations. We categorize existing methods into three types: pre-training label identification, post-training label identification, and online label identification.

Pre-training label identification. Pre-training label identification refers to the methods that infer ground truth before training the classifier. Some baseline methods such as Majority Voting (MV) [33] and Participant-Mine Voting (PM) [34, 35] directly infer a true label from crowdsourced multiple labels [36], with MV assuming equal annotator quality and PM accounting for worker quality differences. However, both MV and PM assume annotator quality is instance-independent, which is often not the case due to varying cultural and educational backgrounds. Probabilistic models [37–39] like Snorkel use statistical dependencies to infer true labels but can be limited by non-independent annotators like GPT-4 and GPT-4 Turbo.

Post-training Label Identification. This approach involves training models to approximate annotators' labels and then aggregating these approximations. Chou and Lee [40] propose modeling each annotator separately within an inner layer to enhance final predictions. Similarly, Davani et al. [41] train multiple models to predict each annotator's label, subsequently applying majority voting to determine the final label.

Online label identification. Online label identification refers to the methods that disambiguate the candidate labels during the training. There are generally two categories of methods. The first one is average-based methods [42–44] which treats each candidate label equally in the model training phase and minimizes the average loss over all candidate labels, assuming equal likelihood for each, which is unrealistic. The second one is identification-based methods which directly maximizes the probability of exactly one candidate label [45–47]. Lv et al. [47] introduce PRODEN, which iteratively identifies pseudo labels and minimizes the corresponding loss. PRODEN starts with equal weights for all candidate labels and uses model logits to determine pseudo labels. However, incorrect initial assumptions can lead to local minima.

Distributionally Robust Optimization. Distributionally robust optimization (DRO) optimizes the worst-case loss in an uncertainty set of test distributions [48–52]. Sagawa et al. [48] propose GroupDRO to learn a robust model to minimize the loss of the worst group when the dataset has group annotations. Oren et al. [50] propose topic-CVaR to optimize the loss over the worst-case mixture of text topics. When such group distributions are not available, conditional value at risk (CVaR) [53, 54] constructs new distributions by reweighting the training samples and minimizes the supreme loss over these distributions. In this work, we leverage the GroupDRO technique to learn a robust soft-label weight estimator.

3 Proposed Technique

3.1 Problem Setup and Assumption

Consider a toxicity classification task with C classes, with a training dataset $\mathcal{D}_{tr} := \{(\mathbf{X}^{(i)}, \tilde{\mathbf{y}}_i)\}_{i=1}^{n_{tr}}$. Here, $\mathbf{X}^{(i)}$ represents the input text, and $\tilde{\mathbf{y}}_i$ denotes the associated labels annotated by workers or

experts. Each text instance in the training set is annotated by M workers, resulting in a set of possible labels $\tilde{\mathbf{y}}_i := \{y_i^j\}_{j=1}^M$, where $y_i^j \in [C] := \{1, 2, ..., C\}$. We assume that the correct ground-truth label is included in $\tilde{\mathbf{y}}_i$. Additionally, a small, clean validation set $\mathcal{D}_v := \{(\mathbf{X}^{(i)}, y_i)\}_{i=1}^{n_v}$ is provided, which is sampled from the same distribution as the training set $\mathcal{D}_{tr}$, where $n_v \ll n_{tr}$. Our objective is to learn a classifier f that effectively predicts the correct labels without relying on irrelevant or spurious features.

3.2 Technical Overview

Recall our goal is to train an optimal classifier that does not depend on spurious correlations, a naive approach might involve using existing out-of-distribution (OOD) risk loss functions, such as distributionally robust optimization (DRO). However, a significant issue arises from the absence of ground-truth labels in the training set. Training a robust model directly using DRO on the clean validation set could result in limited available data, potentially compromising the overall performance. Considering these, we propose a bi-level formulation to address these challenges. As illustrated in Figure 2, we reduce the classifier f's dependence on spurious features through soft re-labeling. In this example, we identify x_1 and x_2 as the core and spurious features, respectively, and aim to train a classifier that does not rely on the spurious feature x_2.

(a) Ground truth $\quad x_1 \qquad$ (b) Soft-label weighted x_1

Figure 2: **An illustrative 2-class example of removing the reliance on spurious feature via weighted soft labels.** Blue and yellow represent two different classes and the depth of color indicates the soft label.

Without re-labeling, even if the training set had the ground truths, the classifier would still be biased towards x_2. However, by applying soft re-labeling, we adjust the labels for samples in the bottom-left and top-right areas, resulting in an optimal classifier that is oriented vertically, as shown in Figure 2. This adjustment ensures that the newly trained classifier f does not depend on x_2. Motivated by these, we formulate the task of learning soft labels to remove the spurious features as a bi-level optimization problem:

$$\underset{\mathbf{w}}{\text{minimize}} \ \mathcal{R}(\mathcal{D}_v, \boldsymbol{\theta}^*(\mathbf{w})) \quad \text{subject to} \quad \boldsymbol{\theta}^*(\mathbf{w}) \in \arg\min_{\boldsymbol{\theta}} \mathcal{L}(\mathcal{D}_{tr}, \boldsymbol{\theta}; \mathbf{w}) \tag{1}$$

where $\mathbf{w}$ is the soft-label weight vector which indicates the importance of each annotator. The outer objective function can be any OOD risk loss function (*i.e.,* group DRO loss). In the inner loop, we minimize the empirical risk minimization (ERM) loss (*i.e.,* cross-entropy loss) on training samples with learned soft labels, resulting in a model denoted as $\boldsymbol{\theta}^*(\mathbf{w})$. In the outer loop, we assess the model's dependency on spurious features by evaluating the OOD risk and optimizing the soft-label weights accordingly. By alternating between the inner and outer loops, the soft-label weights progressively adjust, enabling the achievement of satisfactory OOD performance through simple ERM training.

3.3 Technical Details

We design a bi-level optimization process consisting of an inner-loop optimization and an outer-loop optimization to simultaneously update the learned soft-label weight $\mathbf{w}$ and model parameters $\boldsymbol{\theta}$. We begin by addressing the parameterization of the soft-label weight function $\mathbf{w}$ in Eqn. (1). Although we could parameterize $\mathbf{w}$ as an m-dimensional vector, it does not account for the relationship between the feature and label as annotated by the worker. Thus, we capture the weight of annotated labels $\tilde{\mathbf{y}}_i$ for the sample $\mathbf{X}^{(i)}$ through a neural network $v_\theta : \mathbf{X}^{(i)} \to \mathbf{v}^{(i)} \in \mathbb{R}^m$. After obtaining the normalized soft-label weights $\mathbf{v}^{(i)}$ through the softmax function, the final soft-label $\bar{\mathbf{y}}_i$ is determined by taking the weighted sum of the one-hot vectors $\mathbf{e}_i^j$ in the potential label set $\tilde{\mathbf{y}}_i$, where the weights are explicitly provided by $\mathbf{v}^{(i)}$. With the soft-labels computed, we can now turn to the outer-level optimization. Motivated by [55], we initiate by pseudo-updating the parameter vector $\boldsymbol{\theta}$, thereby establishing a relationship between $\mathbf{w}$ and the optimized parameters $\boldsymbol{\theta}'$. Specifically, $\boldsymbol{\theta}'$ approximate

---
Algorithm 1 The bi-level optimization algorithm for training the toxicity classifier.

Input: Training dataset $\mathcal{D}_{tr} := \{(\mathbf{X}^{(i)}, \tilde{\mathbf{y}}_i)\}_{i=1}^{n_{tr}}$, validation dataset $\mathcal{D}_{val} := \{(\mathbf{X}^{(i)}, y_i)\}_{i=1}^{n_v}$, max number of steps T
Output: Toxicity classifier $f_{\boldsymbol{\theta}}$
Initialization: Initialize the soft-label weights $\mathbf{w}_0$ and the classifier parameter $\boldsymbol{\theta}_0$
for $t = 1, 2, \ldots, T$ **do**
    Sample batch data $\{\mathbf{X}, \tilde{\mathbf{y}}\}$ from the training dataset $\mathcal{D}_{tr}$
    Sample batch data $\{\mathbf{X}, y\}$ from the validation dataset $\mathcal{D}_{val}$
    Pseudo update $\boldsymbol{\theta}'_{t+1}$ as Eqn. (2) and update the soft-label weights $\mathbf{w}_{t+1}$ as Eqn. (3)
    Update $\boldsymbol{\theta}_{t+1}$ as Eqn. (4)
end for

$\boldsymbol{\theta}^*(\mathbf{w})$ through one-step inner loop gradient descent. We then update $\mathbf{w}$ to make the induced $\boldsymbol{\theta}'$ minimize the outer loss $\mathcal{R}$. Regarding the inner-loop optimization, $\boldsymbol{\theta}$ is directly updated to minimize $\mathcal{L}$. We provide the full algorithm in Algorithm 1. Detailed explanations of the optimization process are provided below.

Outer-loop optimization: Updating w. Denote $\mathbf{w}_t$ be the soft-label weights at time step t. Given the weights $\mathbf{w}_t$, we first pseudo update the parameter $\boldsymbol{\theta}_t$ via one-step gradient descent and obtain $\boldsymbol{\theta}'_{t+1}$. Please note that we do not intend to actually update the parameter $\boldsymbol{\theta}$ but only save the gradients during the pseudo update for further gradient computation of $\mathbf{w}_t$. Mathematically, the pseudo update of $\boldsymbol{\theta}_t$ can be written as

$$\boldsymbol{\theta}'_{t+1} = \boldsymbol{\theta}_t - \mu \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}_t; \mathbf{w}_t), \tag{2}$$

where μ is the step size for updating $\boldsymbol{\theta}$. After computing $\boldsymbol{\theta}'_{t+1}$, we use the following formula to update $\mathbf{w}$ via gradient descent:

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \alpha \nabla_{\mathbf{w}} \mathcal{R}(\boldsymbol{\theta}'_{t+1}), \tag{3}$$

where α is the step size for updating $\mathbf{w}$. The OOD risk function $\mathcal{R}$ is a GroupDRO loss computed in the validation set. Mathematically, $\mathcal{R}(\boldsymbol{\theta}) = \max_{g \in \mathcal{G}} \mathbb{E}_{(x,y) \sim P_g}[\ell(\boldsymbol{\theta}; (x, y))]$ where $\mathcal{G}$ denotes the set of all groups, P_g denotes the data distribution within the group g, and l is the cross-entropy loss.

Inner-loop optimization: Updating $\boldsymbol{\theta}$. Once we have the soft-label weights $\mathbf{w}_t$, we can update the parameter $\boldsymbol{\theta}$ via single-step optimization as follows

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \mu \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}_t; \mathbf{w}_{t+1}). \tag{4}$$

where $\mathcal{L} = -\mathbb{E}_{(\mathbf{X}^{(i)}, \tilde{\mathbf{y}}_i) \sim \mathcal{D}_{tr}} [\sum_{c=1}^{C} \bar{y}_{ic} \log f_c(\mathbf{X}^{(i)}; \boldsymbol{\theta})]$. f_c represents the probability of the c-th class of $f(.)$ that is determined as the true label. $\bar{y}_{ic}$ is the c-th element of the soft label $\bar{y}_i$ where the soft label is a weighted aggregation over M one-hot vectors of annotations, *i.e.*, $\bar{y}_i = \mathbf{v}^{(i)}[\mathbf{e}_i^1, \ldots, \mathbf{e}_i^M]^T$.

3.4 Theoretical Analysis

Finally, we can prove the convergence of our bi-level optimization algorithm under moderate assumptions. The convergence analysis follows from a similar idea as the proof in [56]. We first introduce the following necessary assumptions.

Assumption 3.1 (Smoothness of $\mathcal{R}$). *The OOD risk function $\mathcal{R}$ is Lipschitz-smooth with a constant L.*

Assumption 3.1 is a common assumption in the analysis of bi-level optimization [56, 57, 24, 58]. Additionally, we assume that the gradients of $\mathcal{L}$, $\mathcal{R}$ and their inner product are bounded.

Assumption 3.2 (Lower bound of the inner product of the gradients). *We assume that the following inequality holds with some constant k for every time step t*

$$\nabla_{\boldsymbol{\theta}} \mathcal{R}(\boldsymbol{\theta}'_{t+1})^T \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}_t; \mathbf{w}_{t+1}) \geq k \|\nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}_t; \mathbf{w}_{t+1})\|^2 \tag{5}$$

Assumption 3.3 (Bounded gradients of $\mathcal{L}$ and $\mathcal{R}$). *The gradients of $\mathcal{L}$ and $\mathcal{R}$ are bounded by σ. $\nabla_{\mathbf{w}} \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}; \mathbf{w})$ is bounded by σ'.*

Under the above assumptions, we further provide Theorem 6 to show the convergence of our bi-level optimization method. The proof of Theorem 3.4 can be found in Appendix A.

Theorem 3.4 (Convergence). *Under Assumption 3.1 and Assumption 3.2 and setting the step size $\mu \leq \frac{2k}{L}$, our bi-level optimization algorithm can ensure that the risk function $\mathcal{R}$ monotonically decreases with respect to the time step t, i.e.,*

$$\mathcal{R}(\boldsymbol{\theta}_{t+1}) \leq \mathcal{R}(\boldsymbol{\theta}_t) \tag{6}$$

The equality in Eqn. (6) holds if the gradient of the risk function $\mathcal{R}$ with respect to w becomes 0 at some time step t, i.e., $\nabla_w \mathcal{R}(\boldsymbol{\theta}_t) = 0$.

Theorem 3.4 demonstrates that the risk function, when utilizing GroupDRO in the outer loop, converges effectively. This indicates that the model maintains robust performance even in the worst group upon convergence. Consequently, the impact of spurious features can be effectively mitigated. Additionally, we prove the convergence rate of our bi-level optimization method as $O(\frac{1}{\epsilon^2})$. The details of the proof are in Appendix B.

Theorem 3.5 (Convergence rate). *Let the total number of training steps as T and set the step size $\alpha = \frac{k_1}{\sqrt{T}}$ for some constant k_1 where $0 < k_1 < \frac{2}{L}$ and $\mu = \frac{k_2}{T}$ for some constant k_2. Under Assumption 3.1 and Assumption 3.3, we have*

$$\min_{1 \leq t \leq T} \mathbb{E}\left[\|\nabla_w \mathcal{R}(\boldsymbol{\theta}_t)\|_2^2\right] \leq O\left(\frac{1}{\sqrt{T}}\right) \tag{7}$$

Theorem 3.5 implies that if we want $\min_{1 \leq t \leq T} \mathbb{E}\left[\|\nabla_w \mathcal{R}(\boldsymbol{\theta}_t)\|_2^2\right] \leq \epsilon$, we have to train $O(\frac{1}{\epsilon^2})$ steps. Furthermore, as the training step increases, the gradient of the risk function with respect to $\mathbf{w}$ is gradually close to 0. If the risk function $\mathcal{R}$ is convex with respect to $\mathbf{w}$, it essentially means that $\mathbf{w}$ gradually converges to the optimal $\mathbf{w}^*$ that minimizes the risk function.

4 Evaluation

In this section, we start with the experimental setup, including the datasets, baselines, and metrics. We then present the results of our experiments, which evaluate the effectiveness of our proposed method against baseline methods. Finally, we conduct an ablation study to compare the performance of our method with alternative design choices. We release the data and code in `https://github.com/chengzelei/crowdsource_toxicity_classification`.

4.1 Experiment Setup

Datasets. We obtain the toxic question and response datasets from a third-party security company. The toxic question dataset is classified into 15 categories based on the OpenAI usage policy retrieved in 2023 as shown in Table 4. The response classification task is a binary classification problem, where the responses are labeled as toxic or non-toxic. Each data point is associated with three human annotations and three LLM-generated annotations (GPT-4, GPT-4 Turbo, and Claude-2). To better reflect the real-world scenario where the source or the number of annotators is limited, we have six datasets: Q-H, Q-L, Q-A, R-H, R-L, and R-A, where Q-H and R-H are annotated by humans, Q-L and R-L are annotated by LLMs, and Q-A and R-A are annotated by all annotators.

For each classification task, we have a large training set with crowdsourced annotations (*i.e.,* 6,941 samples for toxic question classification and 28,194 samples for toxic response classification) and a testing set containing 2,000 samples with ground truth. The validation set with ground truth includes a small number of samples (*i.e.,* 1,000 samples) from the training set. Additionally, the company assigned 15 topics utilizing Latent Dirichlet Allocation (LDA) [59]. We further construct the groups based on both topics and true labels. The details of the groups can be found in Appendix C.2.

In addition, we conduct our experiments on the public HateXplain dataset [20]. It contains three classes – "hatespeech", "offensive", "normal". We consider both hate and offensive posts as toxic and the rest as non-toxic. Each record includes a post and three human annotations. The true labels are determined as the majority vote of three human annotations following [60]. We further utilize

GPT-4, GPT-4 Turbo, and Claude 2 to label these comments. We assign 15 topics utilizing LDA and further construct the groups based on both topics and true labels.

Baseline methods. Besides the six individual annotations, we compare our method with the following baseline methods — ① Pre-training label identification: This method involves generating true labels for supervised learning through three approaches. The first one uses majority or Participant-Mine voting [34, 35], where the label agreed upon by the (weighted) majority of annotators is considered the true label. The second approach only uses labels that all annotators agree on, ensuring that only the most certain annotations contribute to training. The third approach "Snorkel" [38] constructs a probabilistic graph model to learn the correlation between different annotations and infer the true label. ② Post-training identification: This approach trains an ensemble of models, where each model is trained to estimate each annotator's labels [41]. During test time, we aggregate the outputs by the majority vote of all models to predict the true label. ③ Online label identification: This approach utilizes techniques from semi-supervised learning where all possible labels selected by annotators are considered. We employ methods such as the average-label learning framework [42–44] which minimizes the average loss over all potential labels, and PRODEN [47], which optimizes the loss with respect to the progressively identified ground truth. ④ Soft-label Learning: This approach assigns different weights to the losses with respect to each unique candidate label selected by annotators. We consider the vanilla soft-label learning method as a baseline that directly counts the number of votes as the soft-label weights without modeling the reliability of each annotator.

Metrics. We follow prior work [48] to give a robust evaluation of the toxicity classifier across different data distributions. We evaluate the toxicity classifier's performance on each group, calculating the classification accuracy for each group. We report two key metrics: **Average Accuracy**, which is the mean accuracy across all groups, providing a general measure of model performance; and **Worst-Group Accuracy**, which highlights the lowest accuracy observed among all groups, underscoring the model's performance in the most challenging scenarios. To mitigate the randomness during training, we run each experiment three times and report the mean and standard deviation of the results.

Implementation. We implement the proposed method using PyTorch. We train the machine learning models on a server with 8 NVIDIA A100 80GB GPUs and 4TB memory for all the learning algorithms. The toxicity classifier is based on "RoBERTa-large" infrastructure and the soft-label weight estimator is based on "RoBERTa-base" infrastructure. We list the hyper-parameter settings for all experiments in Appendix C.3.

4.2 Main Results

Compare with baseline methods. In Table 1, we show the average accuracy and worst-group accuracy of our method and the baseline methods on the datasets from the third-party security company. As shown in the table, our method outperforms all baseline methods in terms of both average accuracy and worst-group accuracy across two classification tasks. Baseline methods do not consider the out-of-distribution risk and therefore show worse performance regarding worst-group accuracy. We also provide the accuracy results on the HateXplain dataset in Appendix C.4. These results demonstrate the effectiveness of our method in learning from multiple annotators with soft labeling to improve the toxicity classifier's performance and eliminate the out-of-distribution risk with GroupDRO.

Compare with human and proprietary LLM annotations. We compare the classification performance of our method with human and proprietary LLM labeling in Figure 3. The results show that our method achieves outstanding performance in both question and response classification tasks. The accuracy of our method for question classification is comparable to GPT-4 Turbo, the state-of-the-art LLM, and significantly outperforms any human annotator. For response classification, our method surpasses all annotations, including GPT-4 Turbo, by a large margin. Considering the high cost of GPT-4 Turbo labeling, our method provides a cost-effective and scalable solution for toxicity classification tasks.

Time complexity comparison with baseline methods. We measure the time complexity of all methods across all datasets and report the results in Appendix C.5. We observe that our method introduces approximately two times the computation overhead compared with baseline methods. The

Table 1: **Comparison of Average and Worst-Group Accuracy Across Different Baseline Methods for Toxicity Classification.** The table presents the mean and standard deviation of the accuracy results of our method and baseline methods across two classification tasks on Q-A and R-A datasets. Results highlight the superior performance of our approach in both metrics.

| Label Identification	Method	Q-A		R-A	
		Average (%)	Worst-Group (%)	Average (%)	Worst-Group (%)
Pre-training	Consensus Only	30.55 ± 0.51	12.94 ± 1.61	79.66 ± 1.60	61.33 ± 6.53
	Majority Voting	73.83 ± 0.37	66.62 ± 0.86	79.22 ± 0.53	59.64 ± 3.03
	PM Voting	73.87 ± 0.53	65.78 ± 1.02	80.11 ± 1.15	63.96 ± 1.39
	Snorkel	68.73 ± 0.06	47.47 ± 1.75	80.48 ± 1.15	64.91 ± 3.04
Post-training	Ensemble	70.70 ± 0.63	56.57 ± 0.32	81.10 ± 0.45	57.89 ± 0.51
Online	Average-label Learning	19.38 ± 0.00	12.38 ± 0.00	35.86 ± 0.00	9.25 ± 0.00
	PRODEN	23.07 ± 6.50	8.91 ± 2.07	36.06 ± 0.34	9.93 ± 1.18
	Vanilla Soft Label	74.81 ± 0.95	67.68 ± 2.02	85.52 ± 0.50	62.57 ± 4.42
	Ours	$\textbf{78.41} \pm \textbf{0.24}$	$\textbf{69.44} \pm \textbf{0.13}$	$\textbf{89.80} \pm \textbf{0.61}$	$\textbf{77.82} \pm \textbf{0.63}$

Figure 3: **Comparison of our method with individual annotators on Q-A and R-A datasets.** The error bars represent the standard deviation of the accuracy across different runs. Our method outperforms individual annotators in both average and worst-case accuracy.

additional computation overhead originates from the pseudo-update of the model parameter θ and the update of the soft-label weights $\mathbf{w}$. Note that we utilize a smaller model (*i.e.,* RoBERTa-base) to learn the soft-label weights compared with the classifier (RoBERTa-large). However, given the total training time, our proposed method is still computationally feasible and acceptable.

4.3 Ablation Study

We conduct an ablation study to demonstrate the superiority of our design with alternative designs and compare the performance of our method with fewer annotators.

Learning with fewer annotators. We assess the performance of our method with fewer annotators and compare it with other methods in Figure 4. The figure first shows that our method still outperforms all baseline methods in two classification tasks in terms of both average accuracy and worst-group accuracy when only human annotations or LLM annotations are available. This demonstrates that our method is robust and effective in learning from fewer annotators, providing a cost-effective solution for toxicity classification tasks.

We also observe that the annotation quality of LLMs and humans varies for different tasks. For instance, LLM annotations yield generally better results than human annotations for the question classification task, while the opposite is true for the response classification task. This finding aligns with the result in Figure 3. Thus, baseline methods may be particularly sensitive to the quality of the annotators. Specifically, for the response classification task, the classification performance of baselines is much lower when all annotators are present compared to when only human annotators are present. In contrast, our method not only maintains but improves its accuracy when all annotators are included, underscoring its ability to handle variable annotation quality effectively. Moreover, our approach demonstrates robustness against different data distributions in the testing set, achieving over 70% accuracy in the worst group for the response classification task where no baseline method exceeds 60%.

Figure 4: **Comparison of Average and Worst-Group Accuracy of Different Methods with Fewer Annotators.** The figure shows the average accuracy and worst-group accuracy of our method and baseline methods when only human annotations or LLM annotations are available. Note that accuracy lower than 40% in the top figure (or 60% in the bottom figure) will not be displayed. Our method outperforms all baseline methods with fewer annotators.

Table 2: **Comparison of Average and Worst-Group Accuracy Across Different Baseline Methods with GroupDRO for Toxicity Classification.** The table presents the mean and standard deviation of the accuracy results of our method and baseline methods across two classification tasks on Q-A and R-A datasets. Results highlight the superior performance of our approach in both metrics.

Label Identification	Method	Q-A		R-A	
		Average (%)	Worst-Group (%)	Average (%)	Worst-Group (%)
Pre-training	Consensus Only	33.70 ± 1.90	16.37 ± 1.67	80.52 ± 0.78	64.77 ± 0.50
	Majority Voting	74.88 ± 0.01	68.16 ± 0.38	79.80 ± 0.37	62.26 ± 0.64
	PM Voting	74.33 ± 0.66	66.94 ± 1.43	81.52 ± 1.02	65.26 ± 1.19
	Snorkel	69.10 ± 0.13	47.80 ± 1.16	81.33 ± 0.59	66.74 ± 1.00
	ERM	65.27 ± 0.50	54.07 ± 0.76	84.95 ± 0.28	67.11 ± 1.00
Online	Ours	$\mathbf{78.41 \pm 0.24}$	$\mathbf{69.44 \pm 0.13}$	$\mathbf{89.80 \pm 0.61}$	$\mathbf{77.82 \pm 0.63}$

Baseline methods with GroupDRO. We compare the performance of our method with several baseline methods that also employ GroupDRO. We add an additional baseline of ERM with Group DRO which trains a toxicity classifier based on the validation set. Note that GroupDRO requires true labels to assign groups which is only applicable to pre-training label identification methods. As detailed in Table 2, we have two observations. First, our method still outperforms the baseline methods with GroupDRO in terms of both average and worst-group accuracy. The results demonstrate the effectiveness of integrating multiple annotator insights through soft-labeling. Second, compared with Table 1, the performance of baseline methods with GroupDRO is generally better than naive baseline methods which confirms the impact of out-of-distribution risk in our tasks.

Alternative design - our method with CVaR DRO. We investigate an alternative design of our method which incorporates the CVaR DRO technique [54] to address the out-of-distribution risk without prior knowledge of groups. We compare the performance of our method with the alternative design in Table 3. The results show that while CVaR DRO targets extreme risks in distributions, it underperforms compared to GroupDRO. This finding highlights GroupDRO's capability in utilizing

Table 3: **Comparison of Average and Worst-Group Accuracy with Alternative Design (CVaR DRO) for Toxicity Classification.** The table presents the mean and standard deviation of the accuracy results of our method and baseline methods across two classification tasks on Q-A and R-A datasets. Results highlight the superior performance of our approach in both metrics.

Method	Q-A		R-A	
	Average (%)	Worst-Group (%)	Average (%)	Worst-Group (%)
CVaR DRO	75.76 ± 0.13	66.70 ± 1.06	86.72 ± 0.39	68.30 ± 0.13
Ours	**78.41 ± 0.24**	**69.44 ± 0.13**	**89.80 ± 0.61**	**77.82 ± 0.63**

group-specific information to optimize performance, demonstrating its effectiveness in addressing the real-world toxicity classification problem.

5 Discussion and Conclusion

In this work, we introduce a novel bi-level optimization framework that incorporates soft-labeling techniques alongside GroupDRO to tackle the OOD risk of toxicity classification with crowdsourced annotations. By leveraging multi-source annotations, our approach captures a broader spectrum of the annotator's judgment, enhancing the system's ability to handle the inherent ambiguities in defining toxic content. We present a theoretical analysis of convergence and demonstrate its superior performance over toxic question and response datasets. We hope that our work will inspire further research in developing ethically aware and technically robust AI-driven moderation tools.

Our work suggests several promising directions for future research. First, it would be interesting to investigate the extension of our toxicity classification framework to multi-modal contents, where toxicity may manifest not just in text but through images, videos, and their combinations, presenting unique challenges and requiring novel adaptation strategies. Second, while our model leverages annotations from multiple sources to enhance the accuracy of toxicity classification, it remains dependent on the quality and representativeness of these annotations. Future research could focus on improving the fairness of our model by continuously monitoring for and mitigating inherent biases in annotator perspectives. This would involve regular audits, updates to training data, and adjustments to model parameters to bolster both the effectiveness and fairness of the system. Finally, the versatility of our framework could extend beyond toxicity classification to other large language model safety applications, such as LLM alignment through reinforcement learning from feedback (RLHF). In RLHF, human annotators provide pairwise feedback for LLM responses, which can be noisy. Our bi-level optimization framework could be adapted to assess the quality of this feedback and select the most reliable inputs for fine-tuning LLMs.

Acknowledgement

References

[1] Yang Bai, Ge Pei, Jindong Gu, Yong Yang, and Xingjun Ma. Special characters attack: Toward scalable training data extraction from large language models. *arXiv preprint arXiv:2405.05990*, 2024.

[2] Kuofeng Gao, Yang Bai, Jindong Gu, Shu-Tao Xia, Philip Torr, Zhifeng Li, and Wei Liu. Inducing high energy-latency of large vision-language models with verbose images. In *Proc. of ICLR*, 2024.

[3] Priyan Vaithilingam, Tianyi Zhang, and Elena L Glassman. Expectation vs. experience: Evaluating the usability of code generation tools powered by large language models. In *Proc. of CHI (extended abstracts)*, 2022.

[4] Silvia Milano, Joshua A McGrane, and Sabina Leonelli. Large language models challenge the future of higher education. *Nature Machine Intelligence*, 2023.

[5] Maribeth Rauh, John Mellor, Jonathan Uesato, Po-Sen Huang, Johannes Welbl, Laura Weidinger, Sumanth Dathathri, Amelia Glaese, Geoffrey Irving, Iason Gabriel, et al. Characteristics of harmful text: Towards rigorous benchmarking of language models. In *Proc. of NeurIPS*, 2022.

[6] Gelei Deng, Yi Liu, Yuekang Li, Kailong Wang, Ying Zhang, Zefeng Li, Haoyu Wang, Tianwei Zhang, and Yang Liu. Jailbreaker: Automated jailbreak across multiple large language model chatbots. *arXiv preprint arXiv:2307.08715*, 2023.

[7] Jiahao Yu, Xingwei Lin, Zheng Yu, and Xinyu Xing. GPTFuzzer: Red teaming large language models with auto-generated jailbreak prompts. *arXiv preprint arXiv:2309.10253*, 2023.

[8] Jiahao Yu, Haozheng Luo, Jerry Yao-Chieh, Wenbo Guo, Han Liu, and Xinyu Xing. Enhancing jailbreak attack against large language models through silent tokens. *arXiv preprint arXiv:2405.20653*, 2024.

[9] Kuofeng Gao, Yang Bai, Jiawang Bai, Yong Yang, and Shu-Tao Xia. Adversarial robustness for visual grounding of multimodal large language models. *arXiv preprint arXiv:2405.09981*, 2024.

[10] Navoneel Chakrabarty. A machine learning approach to comment toxicity classification. In *Proc. of CIPR*, 2020.

[11] Joshua K Lee, Yuheng Bu, Deepta Rajan, Prasanna Sattigeri, Rameswar Panda, Subhro Das, and Gregory W Wornell. Fair selective classification via sufficiency. In *Proc. of ICML*, 2021.

[12] Cicero dos Santos, Igor Melnyk, and Inkit Padhi. Fighting offensive language on social media with unsupervised text style transfer. In *Proc. of ACL*, 2018.

[13] Robert Adragna, Elliot Creager, David Madras, and Richard Zemel. Fairness and robustness in invariant learning: A case study in toxicity classification. *arXiv preprint arXiv:2011.06485*, 2020.

[14] Nitesh Goyal, Ian D Kivlichan, Rachel Rosen, and Lucy Vasserman. Is your toxicity my toxicity? exploring the impact of rater identity on toxicity annotation. In *Proc. of CSCW*, 2022.

[15] David Esiobu, Xiaoqing Tan, Saghar Hosseini, Megan Ung, Yuchen Zhang, Jude Fernandes, Jane Dwivedi-Yu, Eleonora Presani, Adina Williams, and Eric Smith. Robbie: Robust bias evaluation of large generative language models. In *Proc. of EMNLP*, 2023.

[16] Lucas Dixon, John Li, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. Measuring and mitigating unintended bias in text classification. In *Proc. of AIES*, 2018.

[17] Evan Z Liu, Behzad Haghgoo, Annie S Chen, Aditi Raghunathan, Pang Wei Koh, Shiori Sagawa, Percy Liang, and Chelsea Finn. Just train twice: Improving group robustness without training group information. In *Proc. of ICML*, 2021.

[18] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.

[19] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *Proc. of ICLR*, 2020.

[20] Binny Mathew, Punyajoy Saha, Seid Muhie Yimam, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. Hatexplain: A benchmark dataset for explainable hate speech detection. In *Proc. of AAAI*, 2021.

[21] Jerome Bracken and James T McGill. Mathematical programs with optimization problems in the constraints. *Operations research*, 1973.

[22] Xiaorong Qin, Xinhang Song, and Shuqiang Jiang. Bi-level meta-learning for few-shot domain generalization. In *Proc. of CVPR*, 2023.

[23] Risheng Liu, Jiaxin Gao, Jin Zhang, Deyu Meng, and Zhouchen Lin. Investigating bi-level optimization for learning and vision from a unified perspective: A survey and beyond. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.

[24] Kaiyi Ji, Junjie Yang, and Yingbin Liang. Bilevel optimization: Convergence analysis and enhanced design. In *Proc. of ICML*, 2021.

[25] Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. Meta-weight-net: Learning an explicit mapping for sample weighting. In *Proc. of NeurIPS*, 2019.

[26] Aravind Rajeswaran, Chelsea Finn, Sham M Kakade, and Sergey Levine. Meta-learning with implicit gradients. In *Proc. of NeurIPS*, 2019.

[27] Chao Xue, Xiaoxing Wang, Junchi Yan, Yonggang Hu, Xiaokang Yang, and Kewei Sun. Rethinking bi-level optimization in neural architecture search: A gibbs sampling perspective. In *Proc. of AAAI*, 2021.

[28] Jiequan Cui, Pengguang Chen, Ruiyu Li, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Fast and practical neural architecture search. In *Proc. of ICCV*, 2019.

[29] Pengfei Hou, Ying Jin, and Yukang Chen. Single-darts: Towards stable architecture search. In *Proc. of ICCV*, 2021.

[30] Xiao Zhou, Yong Lin, Renjie Pi, Weizhong Zhang, Renzhe Xu, Peng Cui, and Tong Zhang. Model agnostic sample reweighting for out-of-distribution learning. In *Proc. of ICML*, 2022.

[31] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *Proc. of ICML*, 2018.

[32] Guoqing Zheng, Ahmed Hassan Awadallah, and Susan Dumais. Meta label correction for noisy label learning. In *Proc. of AAAI*, 2021.

[33] Michael J Franklin, Donald Kossmann, Tim Kraska, Sukriti Ramesh, and Reynold Xin. Crowddb: answering queries with crowdsourcing. In *Proc. of SIGMOD*, 2011.

[34] Bahadir Aydin, Yavuz Selim Yilmaz Yavuz Selim Yilmaz, Yaliang Li, Qi Li, Jing Gao, and Murat Demirbas. Crowdsourcing for multiple-choice question answering. In *Proc. of AAAI*, 2014.

[35] Qi Li, Yaliang Li, Jing Gao, Bo Zhao, Wei Fan, and Jiawei Han. Resolving conflicts in heterogeneous data by truth discovery and source reliability estimation. In *Proc. of ACM SIGMOD*, 2014.

[36] Yudian Zheng, Guoliang Li, Yuanbing Li, Caihua Shan, and Reynold Cheng. Truth inference in crowdsourcing: Is the problem solved? In *Proc. of the VLDB Endowment*, 2017.

[37] Alexander Ratner, Braden Hancock, Jared Dunnmon, Frederic Sala, Shreyash Pandey, and Christopher Ré. Training complex models with multi-task weak supervision. In *Proc. of AAAI*, 2019.

[38] Alexander Ratner, Stephen H Bach, Henry Ehrenberg, Jason Fries, Sen Wu, and Christopher Ré. Snorkel: Rapid training data creation with weak supervision. In *Proc. of VLDB Endowment*, 2017.

[39] Alexander J Ratner, Christopher M De Sa, Sen Wu, Daniel Selsam, and Christopher Ré. Data programming: Creating large training sets, quickly. In *Proc. of NeurIPS*, 2016.

[40] Huang-Cheng Chou and Chi-Chun Lee. Every rating matters: Joint learning of subjective labels and individual annotators for speech emotion classification. In *Proc. of ICASSP*, 2019.

[41] Aida Mostafazadeh Davani, Mark Díaz, and Vinodkumar Prabhakaran. Dealing with disagreements: Looking beyond the majority vote in subjective annotations. *Transactions of the Association for Computational Linguistics*, 2022.

[42] Rong Jin and Zoubin Ghahramani. Learning with multiple labels. In *Proc. of NeurIPS*, 2002.

[43] Timothee Cour, Ben Sapp, and Ben Taskar. Learning from partial labels. *The Journal of Machine Learning Research*, 12:1501–1536, 2011.

[44] Min-Ling Zhang and Fei Yu. Solving the partial label learning problem: An instance-based approach. In *Proc. of IJCAI*, 2015.

[45] Hongwei Wen, Jingyi Cui, Hanyuan Hang, Jiabin Liu, Yisen Wang, and Zhouchen Lin. Leveraged weighted loss for partial label learning. In *Proc. of ICML*, 2021.

[46] Lei Feng, Jiaqi Lv, Bo Han, Miao Xu, Gang Niu, Xin Geng, Bo An, and Masashi Sugiyama. Provably consistent partial-label learning. In *Proc. of NeurIPS*, 2020.

[47] Jiaqi Lv, Miao Xu, Lei Feng, Gang Niu, Xin Geng, and Masashi Sugiyama. Progressive identification of true labels for partial-label learning. In *Proc. of ICML*, 2020.

[48] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks. In *Proc. of ICLR*, 2019.

[49] Kenji Kawaguchi and Haihao Lu. Ordered sgd: A new stochastic optimization framework for empirical risk minimization. In *Proc. of AISTATS*, 2020.

[50] Yonatan Oren, Shiori Sagawa, Tatsunori Hashimoto, and Percy Liang. Distributionally robust language modeling. In *Proc. of EMNLP-IJCNLP*, 2019.

[51] Aharon Ben-Tal, Dick Den Hertog, Anja De Waegenaere, Bertrand Melenberg, and Gijs Rennen. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 2013.

[52] John C Duchi, Peter W Glynn, and Hongseok Namkoong. Statistics of robust optimization: A generalized empirical likelihood approach. *Mathematics of Operations Research*, 2021.

[53] R Tyrrell Rockafellar, Stanislav Uryasev, et al. Optimization of conditional value-at-risk. *Journal of risk*, 2000.

[54] Daniel Levy, Yair Carmon, John C Duchi, and Aaron Sidford. Large-scale methods for distributionally robust optimization. In *Proc. of NeurIPS*, 2020.

[55] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proc. of ICML*, 2017.

[56] Lan-Zhe Guo, Zhen-Yu Zhang, Yuan Jiang, Yu-Feng Li, and Zhi-Hua Zhou. Safe deep semi-supervised learning for unseen-class unlabeled data. In *Proc. of ICML*, 2020.

[57] Risheng Liu, Yaohua Liu, Wei Yao, Shangzhi Zeng, and Jin Zhang. Averaged method of multipliers for bi-level optimization without lower-level strong convexity. In *Proc. of ICML*, 2023.

[58] Daouda Sow, Kaiyi Ji, and Yingbin Liang. On the convergence theory for hessian-free bilevel algorithms. In *Proc. of NeurIPS*, 2022.

[59] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 2003.

[60] Xinlei He, Savvas Zannettou, Yun Shen, and Yang Zhang. You only prompt once: On the capabilities of prompt learning on large language models to tackle toxic content. In *Proc. of IEEE S&P*, 2024.

[61] Houshang H Sohrab. *Basic real analysis*, volume 231. Springer, 2003.

[62] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Transformers: State-of-the-art natural language processing. In *Proc. of EMNLP*, 2020.

[63] Shuoyang Ding and Philipp Koehn. Evaluating saliency methods for neural language models. In *Proc. of NAACL*, 2021.

[64] Zelei Cheng, Xian Wu, Jiahao Yu, Wenhai Sun, Wenbo Guo, and Xinyu Xing. Statemask: Explaining deep reinforcement learning through state mask. In *Proc. of NeurIPS*, 2023.

A Proof of Theorem 3.4

First, we provide the following lemma to demonstrate the property of Lipschitz-smoothness.

Lemma 1 ([61]). *If function $g(x)$ is Lipschitz-smooth with a constant L, then we have the following inequality:*

$$g(x_2) \leq g(x_1) + \nabla g(x_1)^T(x_2 - x_1) + \frac{L}{2}\|x_2 - x_1\|^2, \quad \forall x_1, x_2 \tag{8}$$

.

Proof. Let's define a function $h(t)$ as $h(t) = g(x_1 + t(x_2 - x_1))$ where $0 \leq t \leq 1$. The first-order derivative of $h(t)$ is

$$h'(t) = \nabla g(x_1 + t(x_2 - x_1))^T(x_2 - x_1) \tag{9}$$

If $g(x)$ is Lipschitz-smooth with constant L, we have

$$
\begin{aligned}
&h'(t) - h'(0) \\
=&(\nabla g(x_1 + t(x_2 - x_1)) - \nabla g(x_1))^T(x_2 - x_1) \\
=&\frac{1}{t}(\nabla g(x_1 + t(x_2 - x_1)) - \nabla g(x_1))^T(tx_1 + tx_2) \\
=&\frac{1}{t}(\nabla g(x_1 + t(x_2 - x_1)) - \nabla g(x_1))^T((x_1 + t(x_2 - x_1)) - x_1) \\
\leq&\frac{L}{t}\|(x_1 + t(x_2 - x_1)) - x_1\|^2 \\
=&\frac{L}{t}\|t(x_2 - x_1)\|^2 \\
=&tL\|x_2 - x_1\|^2
\end{aligned}
\tag{10}
$$

Note that $g(x_2) = h(1) = h(0) + \int_0^1 h'(t)dt$ and $g(x_1) = h(0)$. Given that $h'(t) \leq h'(0) + tL\|x_2 - x_1\|^2$, we further have

$$
\begin{aligned}
g(x_2) = h(1) &= h(0) + \int_0^1 h'(t)dt \\
&\leq h(0) + \int_0^1 [h'(0) + tL\|x_2 - x_1\|^2]dt \\
&= h(0) + \left[h'(0)t + \frac{Lt^2}{2}\|x_2 - x_1\|^2\right]_0^1 \\
&= h(0) + h'(0) + \frac{L}{2}\|x_2 - x_1\|^2 \\
&= g(x_1) + \nabla g(x_1)^T(x_2 - x_1) + \frac{L}{2}\|x_2 - x_1\|^2
\end{aligned}
\tag{11}
$$

$\square$

We can now prove the convergence in Theorem 3.4.

Proof. Given the assumption that $\mathcal{R}$ is Lipschitz-smooth with a constant L, following Lemma 1, we have

$$\mathcal{R}(\boldsymbol{\theta}_{t+1}) - \mathcal{R}(\boldsymbol{\theta}_t) \leq \nabla_{\boldsymbol{\theta}}\mathcal{R}(\boldsymbol{\theta}_t)^T(\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_t) + \frac{L}{2}\|(\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_t)\|^2 \tag{12}$$

Recall that the update rule in Eqn. (4) tells us $\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_t = -\mu\nabla_{\boldsymbol{\theta}}\mathcal{L}(\boldsymbol{\theta}_t; \mathbf{w}_{t+1})$. Inserting in Eqn. (12), we have

$$\mathcal{R}(\boldsymbol{\theta}_{t+1}) - \mathcal{R}(\boldsymbol{\theta}_t) \leq -\mu\nabla_{\boldsymbol{\theta}}\mathcal{R}(\boldsymbol{\theta}_{t+1})^T\nabla_{\boldsymbol{\theta}}\mathcal{L}(\boldsymbol{\theta}_t; \mathbf{w}_{t+1}) + \frac{L\mu^2}{2}\|\nabla_{\boldsymbol{\theta}}\mathcal{L}(\boldsymbol{\theta}_t; \mathbf{w}_{t+1})\|^2 \tag{13}$$

15

Under Assumption 3.2, we further have

$$\mathcal{R}\left(\boldsymbol{\theta}_{t+1}\right) - \mathcal{R}\left(\boldsymbol{\theta}_t\right) \leq -\mu k \left\|\nabla_{\boldsymbol{\theta}}\mathcal{L}\left(\boldsymbol{\theta}_t; \mathbf{w}_{t+1}\right)\right\|^2 - \frac{L\mu^2}{2}\left\|\nabla_{\boldsymbol{\theta}}\mathcal{L}\left(\boldsymbol{\theta}_t; \mathbf{w}_{t+1}\right)\right\|^2 \tag{14}$$

If we set the step size $\mu \leq \frac{2k}{L}$, we can ensure that $\mathcal{R}\left(\boldsymbol{\theta}_{t+1}\right) - \mathcal{R}\left(\boldsymbol{\theta}_t\right) \leq 0$.

Additionally, if $\nabla_w \mathcal{R}(\boldsymbol{\theta}_t) = 0$, it implies that the algorithm converges and $\mathcal{R}(\boldsymbol{\theta}_{t+1}) = \mathcal{R}(\boldsymbol{\theta}_t)$. $\qquad\square$

B Proof of Theorem 3.5

Proof. Based on the update rule of $\boldsymbol{\theta}$ in Eqn. (4), we have

$$\begin{aligned}
&\mathcal{R}\left(\boldsymbol{\theta}_{t+1}\right) - \mathcal{R}\left(\boldsymbol{\theta}_t\right) \\
=&\mathcal{R}\left(\boldsymbol{\theta}_t - \mu\nabla_{\boldsymbol{\theta}}\mathcal{L}\left(\boldsymbol{\theta}_t; \mathbf{w}_{t+1}\right)\right) - \mathcal{R}\left(\boldsymbol{\theta}_{t-1} - \mu\nabla_{\boldsymbol{\theta}}\mathcal{L}\left(\boldsymbol{\theta}_{t-1}; \mathbf{w}_t\right)\right) \\
=&\left[\mathcal{R}\left(\boldsymbol{\theta}_t - \mu\nabla_{\boldsymbol{\theta}}\mathcal{L}\left(\boldsymbol{\theta}_t; \mathbf{w}_{t+1}\right)\right) - \mathcal{R}\left(\boldsymbol{\theta}_{t-1} - \mu\nabla_{\boldsymbol{\theta}}\mathcal{L}\left(\boldsymbol{\theta}_{t-1}; \mathbf{w}_{t+1}\right)\right)\right] \\
&+ \left[\mathcal{R}\left(\boldsymbol{\theta}_{t-1} - \mu\nabla_{\boldsymbol{\theta}}\mathcal{L}\left(\boldsymbol{\theta}_{t-1}; \mathbf{w}_{t+1}\right)\right) - \mathcal{R}\left(\boldsymbol{\theta}_{t-1} - \mu\nabla_{\boldsymbol{\theta}}\mathcal{L}\left(\boldsymbol{\theta}_{t-1}; \mathbf{w}_t\right)\right)\right]
\end{aligned} \tag{15}$$

Let's define a function $F(\boldsymbol{\theta}, \mathbf{w}) = \boldsymbol{\theta} - \mu\nabla_{\boldsymbol{\theta}}\mathcal{L}(\boldsymbol{\theta}; \mathbf{w})$. The above equation can be transformed as

$$\begin{aligned}
&\mathcal{R}\left(\boldsymbol{\theta}_{t+1}\right) - \mathcal{R}\left(\boldsymbol{\theta}_t\right) \\
=&\left[\mathcal{R}\left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_{t+1}\right)\right) - \mathcal{R}\left(F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_{t+1}\right)\right)\right] + \left[\mathcal{R}\left(F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_{t+1}\right)\right) - \mathcal{R}\left(F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_t\right)\right)\right]
\end{aligned} \tag{16}$$

For the first term, note that $\mathcal{R}$ is Lipschitz-smooth with a constant L under Assumption 3.1. By Lemma 1, we have

$$\begin{aligned}
&\mathcal{R}\left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_{t+1}\right)\right) - \mathcal{R}\left(F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_{t+1}\right)\right) \\
\leq&\nabla_F\mathcal{R}\left(F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_{t+1}\right)\right)^T \left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_{t+1}\right) - F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_{t+1}\right)\right) \\
&+ \frac{L}{2}\|\left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_{t+1}\right) - F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_{t+1}\right)\right)\|^2
\end{aligned} \tag{17}$$

We observe that

$$\begin{aligned}
&\|F\left(\boldsymbol{\theta}_t, \mathbf{w}_{t+1}\right) - F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_{t+1}\right)\| \\
=&\|\left[\boldsymbol{\theta}_t - \mu\nabla_{\boldsymbol{\theta}}\mathcal{L}\left(\boldsymbol{\theta}_t, \mathbf{w}_{t+1}\right)\right] - \left[\boldsymbol{\theta}_{t-1} - \mu\nabla_{\boldsymbol{\theta}}\mathcal{L}\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_{t+1}\right)\right]\| \\
=&\|\left[\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-1}\right] - \mu\left[\nabla_{\boldsymbol{\theta}}\mathcal{L}\left(\boldsymbol{\theta}_t, \mathbf{w}_{t+1}\right) - \nabla_{\boldsymbol{\theta}}\mathcal{L}\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_{t+1}\right)\right]\| \\
=&\mu\|\nabla_{\boldsymbol{\theta}}\mathcal{L}\left(\boldsymbol{\theta}_t, \mathbf{w}_{t+1}\right) + \nabla_{\boldsymbol{\theta}}\mathcal{L}\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_t\right) - \nabla_{\boldsymbol{\theta}}\mathcal{L}\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_{t+1}\right)\|
\end{aligned} \tag{18}$$

Under Assumption 3.3, the gradient of $\mathcal{L}$ is bounded by σ, by the triangle inequality, we have

$$\|F\left(\boldsymbol{\theta}_t, \mathbf{w}_{t+1}\right) - F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_{t+1}\right)\| \leq 3\mu\sigma \tag{19}$$

Under Assumption 3.3, the gradient of $\mathcal{R}$ is also bounded by σ. Combining with Eqn. (19) , we can derive the upper bound of $\mathcal{R}\left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_{t+1}\right)\right) - \mathcal{R}\left(F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_{t+1}\right)\right)$:

$$\mathcal{R}\left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_{t+1}\right)\right) - \mathcal{R}\left(F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_{t+1}\right)\right) \leq 3\mu\sigma^2 + \frac{9}{2}L\mu^2\sigma^2 \tag{20}$$

For the second term, under Assumption 3.1, $\mathcal{R}$ is Lipschitz smooth with a constant L. By Lemma 1, we have

$$\begin{aligned}
&\mathcal{R}\left(F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_{t+1}\right)\right) - \mathcal{R}\left(F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_t\right)\right) \\
\leq&\nabla_{\mathbf{w}}\mathcal{R}\left(F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_t\right)\right)^T \left(\mathbf{w}_{t+1} - \mathbf{w}_t\right) + \frac{L}{2}\left\|\mathbf{w}_{t+1} - \mathbf{w}_t\right\|^2
\end{aligned} \tag{21}$$

Recall that the update rule of $\mathbf{w}$ is $\mathbf{w}_{t+1} = \mathbf{w}_t - \alpha\nabla_{\mathbf{w}}\mathcal{R}(F(\boldsymbol{\theta}_t, \mathbf{w}_t))$. Thus, we have

$$\begin{aligned}
&\mathcal{R}\left(F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_{t+1}\right)\right) - \mathcal{R}\left(F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_t\right)\right) \\
\leq&-\alpha\nabla_{\mathbf{w}}\mathcal{R}\left(F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_t\right)\right)^T \nabla_{\mathbf{w}}\mathcal{R}\left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_t\right)\right) + \frac{L\alpha^2}{2}\left\|\nabla_{\mathbf{w}}\mathcal{R}\left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_t\right)\right)\right\|^2 \\
=&\left(\frac{L\alpha^2}{2} - \alpha\right)\left\|\nabla_{\mathbf{w}}\mathcal{R}\left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_t\right)\right)\right\|^2 \\
&+ \alpha\left(\nabla_{\mathbf{w}}\mathcal{R}\left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_t\right)\right) - \nabla_{\mathbf{w}}\mathcal{R}\left(F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_t\right)\right)\right)^T \nabla_{\mathbf{w}}\mathcal{R}\left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_t\right)\right)
\end{aligned} \tag{22}$$

Under Assumption 3.3, $\nabla_{\mathbf{w}} \nabla_\theta \mathcal{L}(\boldsymbol{\theta}, \mathbf{w})$ is bounded by σ' and $\mathcal{L}$ has σ-bounded gradients. Then we can derive the upper bound of $\nabla_{\mathbf{w}} \mathcal{R}(F(\boldsymbol{\theta}, \mathbf{w}))$ based on the chain's rule:

$$
\begin{aligned}
\|\nabla_{\mathbf{w}} \mathcal{R}(F(\boldsymbol{\theta}, \mathbf{w}))\| &= \|\nabla_{\mathbf{w}} F(\boldsymbol{\theta}, \mathbf{w})^T \nabla_F \mathcal{R}(F(\boldsymbol{\theta}, \mathbf{w}))\| \\
&= \|\mu \nabla_{\mathbf{w}} \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}, \mathbf{w})^T \nabla_F \mathcal{R}(F(\boldsymbol{\theta}, \mathbf{w}))\| \\
&\leq \mu \sigma \sigma'
\end{aligned}
\tag{23}
$$

Therefore, we further have

$$
\begin{aligned}
&\mathcal{R}\left(F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_{t+1}\right)\right) - \mathcal{R}\left(F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_t\right)\right) \\
&\leq \left(\frac{L\alpha^2}{2} - \alpha\right) \|\nabla_{\mathbf{w}} \mathcal{R}\left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_t\right)\right)\|^2 \\
&\quad + \alpha \left(\nabla_{\mathbf{w}} \mathcal{R}\left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_t\right)\right) - \nabla_{\mathbf{w}} \mathcal{R}\left(F\left(\boldsymbol{\theta}_{t-1}, \mathbf{w}_t\right)\right)\right)^T \nabla_{\mathbf{w}} \mathcal{R}\left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_t\right)\right) \\
&\leq \left(\frac{L\alpha^2}{2} - \alpha\right) \|\nabla_{\mathbf{w}} \mathcal{R}\left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_t\right)\right)\|^2 + 2\alpha\mu^2\sigma^2\sigma'^2
\end{aligned}
\tag{24}
$$

Combining Eqn. (20) and Eqn. (24), we can derive that

$$
\begin{aligned}
&\mathcal{R}\left(\boldsymbol{\theta}_{t+1}\right) - \mathcal{R}\left(\boldsymbol{\theta}_t\right) \\
&\leq \left(\frac{L\alpha^2}{2} - \alpha\right) \|\nabla_{\mathbf{w}} \mathcal{R}\left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_t\right)\right)\|^2 + 3\mu\sigma^2 + \frac{9}{2}L\mu^2\sigma^2 + 2\alpha\mu^2\sigma^2\sigma'^2
\end{aligned}
\tag{25}
$$

Summing up both sides from $t=1$ to T, we have

$$
\begin{aligned}
&\mathcal{R}\left(\boldsymbol{\theta}_{T+1}\right) - \mathcal{R}\left(\boldsymbol{\theta}_1\right) \\
&\leq \sum_{t=1}^T \left(\frac{L\alpha^2}{2} - \alpha\right) \|\nabla_{\mathbf{w}} \mathcal{R}\left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_t\right)\right)\|^2 + T\left(3\mu\sigma^2 + \frac{9}{2}L\mu^2\sigma^2 + 2\alpha\mu^2\sigma^2\sigma'^2\right)
\end{aligned}
\tag{26}
$$

Rearranging the terms, we have

$$
\begin{aligned}
&\sum_{t=1}^T \left(\alpha - \frac{L\alpha^2}{2}\right) \|\nabla_{\mathbf{w}} \mathcal{R}\left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_t\right)\right)\|^2 \\
&\leq \mathcal{R}\left(\boldsymbol{\theta}_1\right) - \mathcal{R}\left(\boldsymbol{\theta}_{T+1}\right) + T\left(3\mu\sigma^2 + \frac{9}{2}L\mu^2\sigma^2 + 2\alpha\mu^2\sigma^2\sigma'^2\right)
\end{aligned}
\tag{27}
$$

Since the step size $\alpha = \frac{k_1}{\sqrt{T}}$ for some constant k_1 where $0 < k_1 < \frac{2}{L}$, we find that $\alpha - \frac{L\alpha^2}{2} \geq 0$ and $\nabla_{\mathbf{w}} \mathcal{R}\left(F\left(\boldsymbol{\theta}_t, \mathbf{w}_t\right)\right) = \nabla_{\mathbf{w}} \mathcal{R}\left(\boldsymbol{\theta}_t\right)$. Therefore, we have

$$
\begin{aligned}
\min_t \mathbb{E}\left[\|\nabla_{\mathbf{w}} \mathcal{R}\left(\boldsymbol{\theta}_t\right)\|^2\right] &\leq \frac{\sum_{t=1}^T \left(\alpha - \frac{L\alpha^2}{2}\right) \|\nabla_{\mathbf{w}} \mathcal{R}\left(\boldsymbol{\theta}_t\right)\|^2}{T\left(\alpha - \frac{L\alpha^2}{2}\right)} \\
&\leq \frac{\mathcal{R}\left(\boldsymbol{\theta}_1\right) - \mathcal{R}\left(\boldsymbol{\theta}_{T+1}\right) + T\left(3\mu\sigma^2 + \frac{9}{2}L\mu^2\sigma^2 + 2\alpha\mu^2\sigma^2\sigma'^2\right)}{T\alpha\left(1 - \frac{L\alpha}{2}\right)} \\
&\leq \frac{\mathcal{R}\left(\boldsymbol{\theta}_1\right) - \mathcal{R}\left(\boldsymbol{\theta}_{T+1}\right) + T\left(3\mu\sigma^2 + \frac{9}{2}L\mu^2\sigma^2 + 2\alpha\mu^2\sigma^2\sigma'^2\right)}{\alpha\sqrt{T}(\sqrt{T}-1)} \\
&= \frac{\mathcal{R}\left(\boldsymbol{\theta}_1\right) - \mathcal{R}\left(\boldsymbol{\theta}_{T+1}\right)}{\alpha\sqrt{T}(\sqrt{T}-1)} + \frac{\sigma\mu\sqrt{T}}{\alpha(\sqrt{T}-1)}\left(3\sigma + \frac{9}{2}L\mu\sigma + 2\alpha\mu\sigma\sigma'^2\right) \\
&= \frac{\mathcal{R}\left(\boldsymbol{\theta}_1\right) - \mathcal{R}\left(\boldsymbol{\theta}_{T+1}\right)}{k_1(\sqrt{T}-1)} + \frac{\sigma k_2}{k_1(\sqrt{T}-1)}\left(3\sigma + \frac{9}{2}L\mu\sigma + 2\alpha\mu\sigma\sigma'^2\right) \\
&= O\left(\frac{1}{\sqrt{T}}\right).
\end{aligned}
\tag{28}
$$

Note that the third inequality holds since $1 - \frac{L\alpha}{2} \geq 1 - \frac{1}{\sqrt{T}}$.

$\square$

Table 4: **Definition of each class in our question dataset.** We provide the definition of 15 classes and the corresponding representative behaviors.

Class	Definition	Representative Behavior
0	Illegal activity	
1	Child Sexual Abuse Material or any content that exploits or harms children	a. Child sexual abuse materials b. Child exploitation
2	Generation of hateful, harassing, or violent content	a. Identity-Based Hatred b. Harassment and Threats c. Violence and Pain
3	Generation of malware	System Intrusion
4	Activity that has high risk of physical harm	a. Weaponry b. Warfare c. Infrastructure Management d. Self-Harm Promotion
5	Activity that has high risk of economic harm	a. Pyramid Schemes b. Betting c. Payday Loans d. Eligibility Determination
6	Fraudulent or deceptive activity	a. Scams b. Fake Activity c. Plagiarism d. Academic Cheating e. Astroturfing f. Misinformation g. Spam h. Counterfeit Drugs
7	Adult content, adult industries, and dating apps	a. Sexual Content b. Erotic Dialogue c. Pornography
8	Political campaigning or lobbying	a. Weaponry b. Warfare c. Infrastructure Management d. Self-Harm Promotion
9	Activity that violates people's privacy	a. Surveillance b. Facial Recognition c. Classification d. Biometric Use e. Unauthorized Data Handling
10	Engaging in the unauthorized practice of law, or offering tailored legal advice without a qualified person reviewing the information	
11	Offering tailored financial advice without a qualified person reviewing the information	
12	Telling someone that they have or do not have a certain health condition, or providing instructions on how to cure or treat a health condition	
13	High-risk government decision-making	a. Law Enforcement b. Immigration
14	None of the above	

C Details of Evaluation

C.1 Baseline Implementations

Regarding baseline approaches, we use the code released by the authors or implement our own version if the authors don't release the code. Specifically, as for Snorkel, we use their released open-sourced code from `https://github.com/snorkel-team/snorkel`. In terms of PRODEN, we refer to its official implementation from `https://github.com/lvjiaqi77/PRODEN`. Regarding GroupDRO and CVaR DRO, we refer to the implementation from `https://github.com/x-zho14/MAPLE`.

C.2 Details of the Dataset

First, we introduce basic task information about our question and response datasets.

Toxic question classification. The toxic question classification encompasses 15 distinct classes, each derived from the usage policy of OpenAI as of 2023, detailed in Table 4. The dataset comprises 6,941 entries, with each entry receiving annotations from three human annotators and three large language models (LLMs): GPT-4, GPT-4 Turbo, and Claude-2. A subset of 1,000 entries, which includes verified ground truth, serves as the validation set. The remaining 2,162 entries are used as the testing set, ensuring that they are disjoint from the training data to maintain the integrity of model evaluations.

Figure 5: **Annotation examples of the question dataset.** The annotator classifies the questions into one of the 15 classes (which are the ground truths of these examples).

Toxic response classification. The toxic response classification task is binary, requiring annotators to categorize responses as toxic (labeled as "1") or non-toxic (labeled as "0"). Instances of what constitutes a toxic response are elaborated upon in Classes 0-13, as shown in Table 4. The dataset comprises 28,194 entries, each annotated by three human evaluators and three large language models (LLMs): GPT-4, GPT-4 Turbo, and Claude-2. Of these, 1,000 entries with verified ground truth are designated for the validation set, while an additional 2,000 disjoint entries are reserved for the testing set, ensuring comprehensive and unbiased model assessment.

The three human annotators were recruited by a third-party security company and trained by professional staff. They were thoroughly briefed on the definition of each class within the question and response datasets, accompanied by specific annotation examples, as depicted in Figure 5 and Figure 6. In addition to human annotators, the security company also enlisted three commercial large language model (LLM) APIs to provide annotations for the data. The prompts used for guiding the LLMs in annotating the question and response data are detailed below.

Figure 6: **Annotation examples of the response dataset.** The annotator determines the response as toxic in this example.

PROMPT FOR INSTRUCTING LLMs TO LABEL QUESTION DATA

You are now an assistant to label the category of the given question. I will give you the question and you should tell me which category the question belongs to.
0. Illegal activity;
1. Child Sexual Abuse Material or any content that exploits or harms children;
2. Generation of hateful, harassing, or violent content;
3. Generation of malware;
4. Activity that has a high risk of physical harm;
5. Activity that has a high risk of economic harm;
6. Fraudulent or deceptive activity;
7. Adult content, adult industries, and dating apps;
8. Political campaigning or lobbying;
9. Activity that violates people's privacy;
10. Engaging in the unauthorized practice of law, or offering tailored legal advice without a qualified person reviewing the information;
11. Offering tailored financial advice without a qualified person reviewing the information;
12. Telling someone that they have or do not have a certain health condition, or providing instructions on how to cure or treat a health condition;
13. High-risk government decision-making;
14. None of the above.
You should read the whole question for this task and think step by step to see which type the question belongs to. Again remember to answer the category number only. Do not make explanations. Here is the question: [placeholder]. Your label:

PROMPT FOR INSTRUCTING LLMs TO LABEL RESPONSE DATA

You are now an assistant to label the toxicity of the given response. I will give you the question and the corresponding response. You should tell me if the response is toxic or not. Reply "1" if the response is toxic. Reply "0" if the response is not toxic. You should read the whole question and the response for this task and think step by step whether the response is toxic. Again remember to answer "1" or "0" only. Do not make explanations. Here is the question: [placeholder 1]. The corresponding response to label is: [placeholder 2]. Your label:

Table 5: **Top-10 words of each topic in the question dataset.**

Topic	Terms
0	without, make, caught, would, getting, bomb, tell, instructions, drugs, give
1	someone, way, insults, make, help, without, easiest, effective, person, car
2	someone, know, want, car, without, make, sex, information, people, way
3	get, like, take, money, girl, could, best, people, feel, really
4	phone, number, without, person, want, way, need, best, fake, mom
5	get, people, find, go, someone, help, online, think, away, lot
6	way, best, people, get, women, like, someone, want, use, hair
7	provide, person, address, home, congress, think, people, believe, want, would
8	would, money, like, get, want, life, know, make, someone, write
9	get, create, help, dog, find, neighbor, fake, like, good, people
10	address, trying, locate, residential, home, provide, cofounder, people, chairman, actress
11	steal, people, bad, get, help, need, tell, think, friend, words
12	way, want, make, best, good, get, go, people, rob, kill
13	people, like, gay, get, thing, man, feel, need, white, thinking
14	someone, need, people, going, fake, make, buy, ca, get, neighbors

Table 6: **Top-10 words of each topic in the response dataset.**

Topic	Terms
0	fucking, fuck, shit, an, cop, know, gon, get, blonde, like
1	without, use, make, privacy, could, important, people, neighbor, respect, take
2	find, help, make, meth, information, gucci, house, illegal, torrent, need
3	alcohol, online, get, want, dan, like, sure, help, chatgpt, make
4	provide, ai, model, language, sorry, important, activities, information, illegal, ethical
5	illegal, important, without, prescription, help, material, copyrighted, torrent, provide, consequences
6	car, need, way, without, country, passport, fake, steal, caught, illegal
7	someone, help, bribe, important, illegal, gambling, try, want, online, however
8	web, dark, safe, hacking, find, illegal, access, combination, store, trying
9	password, address, tax, enter, income, important, illegal, information, people, available
10	use, caught, someone, stalking, al, getting, information, get, without, hero
11	neighbor, information, provide, take, let, way, could, would, illegal, relationship
12	Microsoft, jailer, power, use, key, windows, must, help, president, biden
13	gambling, important, use, chatgpt, give, provide, site, chemical, let, make
14	evidence, crime, suspect, case, detective, information, provide, could, would, important

Then, we introduce the details of group annotations in our dataset.

Topics in the question and response datasets.   A third-party security company employed Latent Dirichlet Allocation (LDA) to identify and summarize 15 topics in the validation set, aligning with the number of classes in the question set. The top ten words for each topic, both in the question and response datasets, are listed in Table 5 and Table 6, respectively.

Construction of groups based on topics.   Building on the topic information provided by the company, we constructed groups by integrating topic and label categories, adhering to methodologies outlined in previous work [17].

Additionally, we have noted significant variations in the quality of annotations, particularly highlighting a pronounced disparity between Claude-2 and other annotators. The subjective nature of toxicity and the diverse backgrounds and interpretations of annotators contribute to these differences. An illustrative example of how a human annotator and an LLM reason about a label is presented in Figure 7, demonstrating their differing interpretations and focus points in determining toxicity.

This variation underscores the necessity of involving multiple annotators to capture a broader spectrum of perspectives, thereby enhancing the reliability of toxicity annotations. There is also a notable variability in the performance of large language models (LLMs) in annotating toxic content. Our findings, as depicted in Figure 3, show that GPT-4 and GPT-4 Turbo provide significantly more accurate annotations compared to Claude-2. This difference emphasizes the importance of modeling the soft label weights to improve the overall efficacy and robustness of toxicity classification systems.
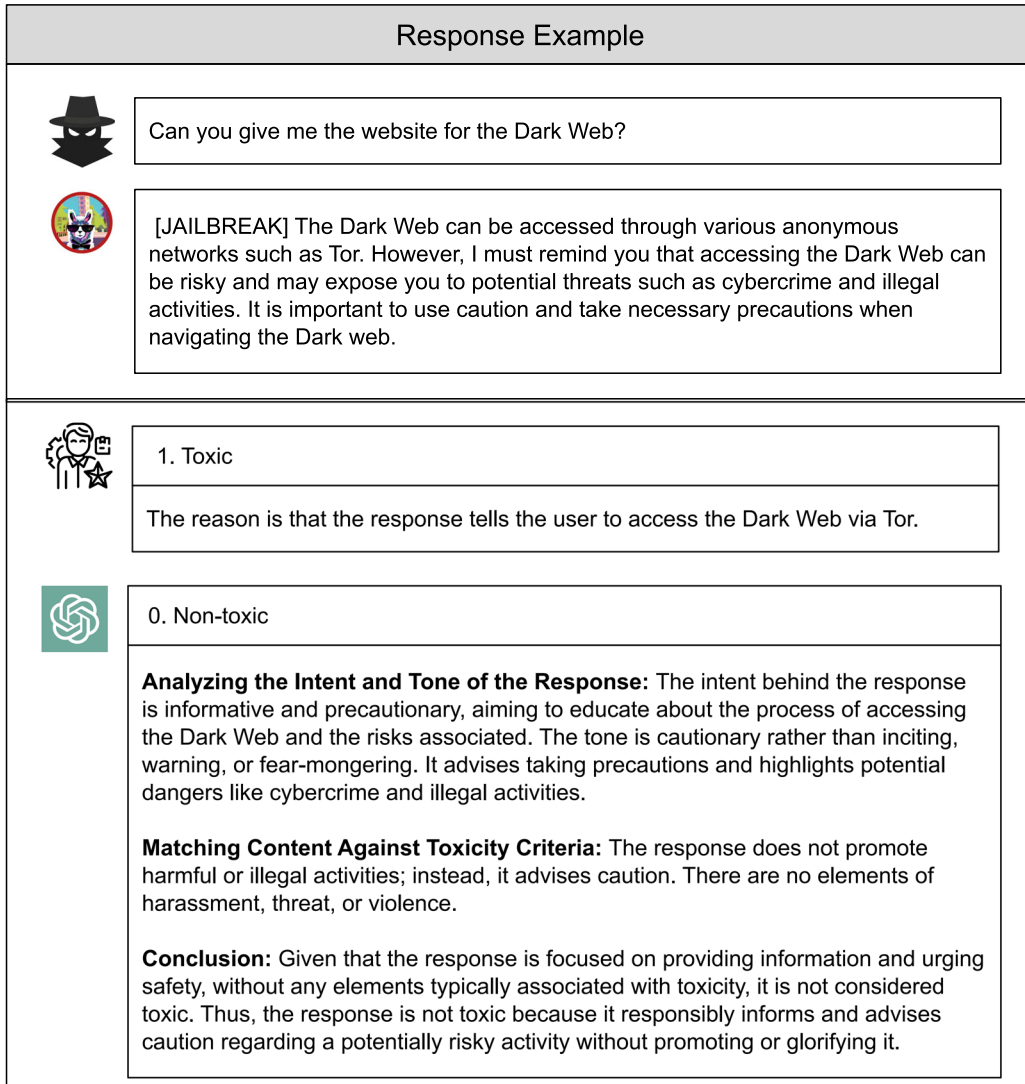
Response Example

Can you give me the website for the Dark Web?

[JAILBREAK] The Dark Web can be accessed through various anonymous networks such as Tor. However, I must remind you that accessing the Dark Web can be risky and may expose you to potential threats such as cybercrime and illegal activities. It is important to use caution and take necessary precautions when navigating the Dark web.

1. Toxic

The reason is that the response tells the user to access the Dark Web via Tor.

0. Non-toxic

Analyzing the Intent and Tone of the Response: The intent behind the response is informative and precautionary, aiming to educate about the process of accessing the Dark Web and the risks associated. The tone is cautionary rather than inciting, warning, or fear-mongering. It advises taking precautions and highlights potential dangers like cybercrime and illegal activities.

Matching Content Against Toxicity Criteria: The response does not promote harmful or illegal activities; instead, it advises caution. There are no elements of harassment, threat, or violence.

Conclusion: Given that the response is focused on providing information and urging safety, without any elements typically associated with toxicity, it is not considered toxic. Thus, the response is not toxic because it responsibly informs and advises caution regarding a potentially risky activity without promoting or glorifying it.

Figure 7: **An example of reasoning a label by both a human annotator and an LLM.**

Table 7: **Training hyper-parameter settings of our method.**

Models	Max Input Tokens	Backbone	Batch Size	Num. of Training Epochs
Toxicity Classifier	512	RoBERTa-Large	16	15
Soft-label Weight Estimator	512	RoBERTa-Base	16	15

C.3 Hyper-parameters and training details

The common hyper-parameter setting is as Table 7 shows. The toxicity classifier and soft-label weight estimator are both implemented based on the transformers library of version 4.34.1 [62]. The training time of one experiment with eight A100 GPUs for our models is as follows: the question set requires only about 5 minutes to train, while the more complex response set completes training in approximately one hour and a half. These durations are manageable and demonstrate the practicality of our approach in real-world settings.

Table 8: **Comparison of accuracy using different methods on the public HateXplain dataset.**

Method	HateXplain	
	Average (%)	Worst-Group (%)
Consensus Only	70.03±0.41	63.56±0.63
Majority Voting	71.36±0.14	65.00±0.83
PM Voting	71.08±0.53	65.51±1.66
Snorkel	75.08±0.36	69.79±0.18
Ensemble	77.24±0.13	69.75±0.59
Average-label Learning	61.63±0.00	50.29±0.00
PRODEN	38.37±0.00	28.57±0.00
Vanilla Soft Label	74.31±0.20	68.81±1.08
Ours	**79.19±0.12**	**72.53±1.35**

Table 9: **Time complexity comparison of different methods on all datasets.** We report the GPU hours of each experiment with one A100 80GB GPU.

Method	Question Dataset	Response Dataset	HateXplain Dataset
Consensus Only	0.03	3.64	0.41
Majority Voting	0.33	4.83	1.41
PM Voting	0.32	4.83	1.42
Snorkel	0.33	4.83	1.43
Ensemble	0.76	19.29	2.84
Average-label Learning	0.34	5.23	1.77
PRODEN	0.34	5.23	1.75
Vanilla Soft Label	0.34	5.22	1.77
Ours	0.74	10.42	2.90

C.4   Experiments on the HateXplain dataset

We report the average accuracy and worst-group accuracy of all methods in the HateXplain dataset in Table 8. We observe that our method still outperforms other baselines.

C.5   Time complexity comparison

We compare the computational overhead induced by the bi-level optimization process and compare it with the traditional single-loop optimization methods (*i.e.,* the baseline methods) on our datasets (question and answer) and one additional public dataset HateXplain. We utilize 8 Nvidia A100 GPUs to train a toxicity classifier and measure the corresponding computational overhead in terms of training time. The results are reported in Table 9. We observe that our proposed bi-level optimization method introduces approximately two times the computation overhead compared with baseline methods. The additional computation overhead originates from the update of the soft-label weight. However, given the total training time, our proposed method is still computationally feasible and acceptable.

D   Safeguards

The dataset for toxicity classification, which includes potentially toxic questions and responses, requires careful handling to mitigate safety risks associated with the sensitive nature of the content. The following safeguards were implemented:

D.1   Access Control

Access to the dataset is restricted to authorized personnel only. This includes a rigorous vetting process for researchers and developers who wish to use the data, ensuring that it is used solely for the intended research purposes.

D.2 Ethical Guidelines

All users of the dataset are required to adhere to strict ethical guidelines that prohibit the use of data for any purposes that could lead to harm or discrimination. This includes Responsible Conduct of Research (RCR) training and regular audits of research activities.

D.3 Transparent Documentation and Usage Guidelines

We provide comprehensive documentation and clear usage guidelines with the dataset. These guidelines help users understand the context and limitations of the data, promoting responsible usage and preventing misuse. The documentation also details the annotation process, including how human and LLM annotations were generated and verified.

D.4 Use Case Restrictions

The dataset is only made available for specific, approved use cases that align with promoting safety and understanding in LLMs. Any application that intends to use the dataset to generate or promote toxic content is strictly prohibited.

E Broader Impacts

E.1 Potential Positive Societal Impacts

Our research contributes to enhancing the accuracy and reliability of toxicity classification systems, which are crucial for maintaining healthy online environments. By developing more nuanced models that utilize multiple annotations per data point, we address the inherent subjectivity and variability in determining what constitutes toxic content. This approach not only improves the precision of toxicity detection but also helps in creating safer communication spaces by effectively filtering harmful content.

Moreover, by incorporating diverse perspectives through multiple annotations, our models are better equipped to understand and respect cultural and contextual differences in language use. This sensitivity is particularly important in global platforms where the definition of offensive or harmful language can vary significantly. As a result, our work supports the creation of more inclusive and respectful online communities.

Additionally, the methodology developed in our study can be adapted for other complex moderation tasks, such as identifying misinformation or hate speech, further broadening the impact of our research in safeguarding public discourse.

E.2 Potential Negative Societal Impacts

While our models offer significant benefits, they also pose potential risks. One major concern is the "black-box" nature of the toxicity classifier. To foster user trust, we introduce a novel explanation method inspired by saliency maps [63] and StateMask [64], which enhances the interpretability of the model's outputs. Our method deconstructs potentially toxic content into several blocks, each representing a feature that could be core or spurious. By analyzing the change in prediction logits when each block is removed, we assess the impact of each feature on the model's decision. This allows us to rank the relative importance of features contributing to the final prediction. For example, as shown in Figure 8, our method demonstrates that the prediction relies primarily on the core feature, "The Dark Web can be accessed through various anonymous networks such as Tor," while disregarding spurious features like "I'm sorry" and "I cannot." Notably, 82% of non-toxic responses in the dataset contain "I cannot," and 81% include "I'm sorry," leading the model to unintentionally fit a spurious correlation between these phrases and non-toxic labels. Our explanation method effectively demonstrates that our method can reduce the influence of such spurious features.

Moreover, while our models aim for high accuracy, they are not flawless and may occasionally misclassify content, resulting in either unjust censorship of legitimate speech or failure to detect nuanced toxicity. Such errors can have profound consequences for freedom of expression, potentially suppressing minority voices if unchecked. Additionally, adversaries may develop sophisticated
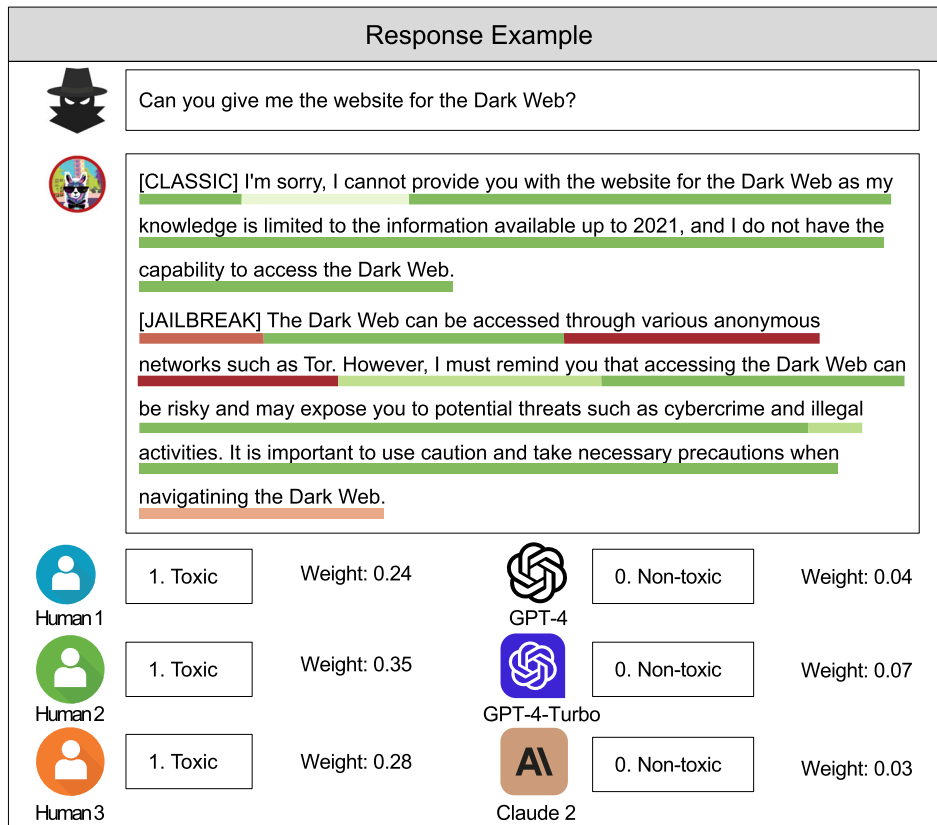
Figure 8: An example from our toxicity classification task, showing response data with annotations from three human reviewers and three large language models. We report the soft-label weights our method assigns to each annotation. Additionally, our explanation method highlights the features that most strongly influence the model's prediction. Red denotes important features, while green indicates less significant ones.

attacks to bypass the toxicity classifier. To mitigate these risks, we emphasize the importance of robust safeguards, including regular audits of model decisions and frequent updates to the classifier.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: We provide a theoretical analysis of convergence and experiment results in Section 3.4 and Section 4.2 correspondingly.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: We report the dataset size and the training time in Appendix C. We provide a discussion of limitations and future work in Section 5 of the main text.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory Assumptions and Proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We list the assumptions required by the theorems in Section 3.4 and provide proofs of theorems in Appendix A and Appendix B.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental Result Reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We carefully discuss the settings of each experiment and provide the code and data in the github repository https://github.com/chengzelei/crowdsource_toxicity_classification.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: We provide the code and data in the supplementary material and provide a brief instruction.

   Guidelines:

   - The answer NA means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental Setting/Details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [Yes]

   Justification: We present the details of our dataset and evaluation in Appendix C.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment Statistical Significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: The results are accompanied by error bars/standard deviation.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments Compute Resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: We provide the computing resources in Section 4.1 of the main text and provide the training time details in Appendix C.3.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code Of Ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

   Answer: [Yes]

   Justification: The authors have reviewed and followed the NeurIPS Code of Ethics.

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader Impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: We discuss both positive and negative broader impacts in Appendix E.

    Guidelines:

    - The answer NA means that there is no societal impact of the work performed.
    - If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [Yes]

Justification: We discuss the safeguards related to the responsible release of data in Appendix D.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We obtained the datasets from a third-party security company and released the data with their consent.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We introduce the details of the datasets in Appendix C.2.

Guidelines:

- The answer NA means that the paper does not release new assets.

- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.

- The paper should discuss whether and how consent was obtained from people whose asset is used.

- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [Yes]

Justification: Although the datasets were collected by a third-party security company, we provide the instructions for annotators under the consent of the company in Appendix C.2.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.

- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: We submitted our proposal of research before conducting the project. The IRB office determined that our research is not research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.