

Insights from the Inverse: Reconstructing LLM Training Goals Through Inverse Reinforcement Learning

Jared Joselowitz, Ritam Majumdar, Arjun Jagota, Matthieu Bou, Nyal Patel
Imperial College London

Satyapriya Krishna
Harvard University

Sonali Parbhoo*
Imperial College London
s.parbhoo@imperial.ac.uk

WARNING: This paper contains model outputs that may be considered offensive.

Abstract

Large language models (LLMs) trained with Reinforcement Learning from Human Feedback (RLHF) have demonstrated remarkable capabilities, but their underlying reward functions and decision-making processes remain opaque. This paper introduces a novel approach to interpreting LLMs by applying inverse reinforcement learning (IRL) to recover their implicit reward functions. We conduct experiments on toxicity-aligned LLMs of varying sizes, extracting reward models that achieve up to 85% accuracy in predicting human preferences. Our analysis reveals key insights into the non-identifiability of reward functions, the relationship between model size and interpretability, and potential pitfalls in the RLHF process. We demonstrate that IRL-derived reward models can be used to fine-tune new LLMs, resulting in comparable or improved performance on toxicity benchmarks. This work provides a new lens for understanding and improving LLM alignment, with implications for the responsible development and deployment of these powerful systems.¹

1 Introduction

Large language models (LLMs) have achieved remarkable success in natural language processing, powering applications like conversational AI, translation, and content moderation. A key enabler of these advances is Reinforcement Learning from Human Feedback (RLHF) (Casper et al., 2023), which aligns model outputs with human preferences via reward signals. However, the reward functions learned during RLHF remain opaque, raising concerns about interpretability and safety, particularly in high-stakes domains (Liao & Vaughan, 2023; Liu et al., 2023).

Inverse Reinforcement Learning (IRL) is widely used in robotics and control theory to infer latent reward structures from observed behavior (Ng & Russell, 2000). By treating observed actions as expert demonstrations, IRL techniques aim to reverse-engineer the hidden reward function that an agent is implicitly optimizing. In this work, we propose a novel application of IRL to LLMs trained via RLHF. We posit that if the outputs of an RLHF-trained LLM can be interpreted as demonstrations from an “expert” policy, then IRL methods—particularly those based on maximum margin formulations—can be employed to recover the hidden reward functions that guided the training process. This approach provides insights into decision-making processes and enhances model auditing.

¹Code for our paper can be found at https://github.com/ai4ai-lab/irl_for_llms

Our work is motivated by key observations: The opacity of RLHF-derived rewards can enable reward hacking (Skalse et al., 2022)—where models exploit spurious correlations rather than genuinely aligning with human values. While prior interpretability studies (Yuan et al., 2023; Dong et al., 2023; Zhao et al., 2023; Azar et al., 2024) have focused on dissecting model architectures, they do not explicitly reveal the incentives shaping LLM behavior. IRL, in contrast, offers a principled framework for uncovering these hidden reward structures.

We employ Maximum Margin IRL to learn a reward model $\hat{R} : \mathcal{O} \rightarrow \mathbb{R}$ that distinguishes between preferred (non-toxic) and suboptimal (toxic) responses. By parameterizing \hat{R} using hidden representations from a base LLM and optimizing it with an asymmetric margin loss, we enforce clear reward separation while addressing non-identifiability challenges. Our experiments focus on toxicity reduction, evaluating Pythia models (70M and 410M parameters) (Biderman et al., 2023) on Jigsaw Toxicity and RealToxicityPrompts benchmarks. Results show that IRL-extracted reward functions closely align with human judgments. We quantitatively evaluate the recovered rewards using metrics such as classification accuracy, recall, ranking correlation, and separation metrics, and further explore the robustness of the inferred reward functions under various perturbations and noisy conditions.

Our key contributions are threefold:

1. **IRL Framework for LLMs:** We introduce an IRL-based method to extract latent reward functions from RLHF-trained models using a maximum margin approach.
2. **Comprehensive Empirical Analysis:** We demonstrate that IRL-recovered rewards capture RLHF objectives while revealing vulnerabilities like non-identifiability and reward sensitivity.
3. **Implications for Model Auditing:** By exposing underlying reward structures, our method provides a diagnostic tool for auditing LLM safety and alignment.

By exposing the hidden incentives that drive LLM behavior, our approach not only enhances our theoretical understanding of RLHF but also paves the way for practical interventions that can mitigate model deployment risk. The rest of the paper details the technical aspects of our IRL framework, present extensive empirical results, and discusses future directions for integrating IRL into the broader landscape of model auditing and safety research.

2 Methodology

We propose a framework for recovering the reward function used to fine-tune a LLM via RLHF by leveraging IRL. Our pipeline consists of four key steps: (i) data curation and processing, (ii) training a groundtruth reward model, (iii) fine-tuning LLMs with RLHF using the groundtruth reward model, and (iv) applying IRL to approximate the underlying reward function from the fine-tuned LLM. Finally, we evaluate the extracted reward function \hat{R} by comparing it to the true reward model R^* to assess what properties of R^* are captured by \hat{R} . We describe each stage below.

2.1 Data Processing.

Let $\mathcal{D} = \{(c_i, y_i)\}_{i=1}^N$ be a dataset where c_i represents a comment and $y_i \in \mathbb{R}$ denotes the label. In our work, we focus on the task of toxicity (lower values correspond to non-toxic, while higher values to toxic). We construct a balanced dataset $\mathcal{D}_{\text{bal}} \subset \mathcal{D}$ containing equal numbers of toxic and non-toxic samples. Each data point is further split into prompt-output pairs (p_i, o_i) for training and evaluation purposes. The resulting \mathcal{D}_{bal} dataset forms the basis for both RLHF fine-tuning, IRL reward extraction and evaluation.

2.2 Extracting the Ground Truth Reward Model R^* and Fine Tuning LLM using R^*

Extracting Ground Truth R^* . An effective reward function is fundamental to the RLHF methodology, acting as an automated substitute for human input. We assume the true reward function R^* is unknown at training time, as is the case in most practical scenarios,

but can be parameterised by the function $f_\theta : \mathcal{O} \rightarrow \mathbb{R}$, where \mathcal{O} is the space of text outputs. That is, R^* can distinguish between toxic and non-toxic comments (0 = non-toxic, 1 = toxic). Formally, the reward function is: $R^*(o) = -f_\theta(o)$, where $o \in \mathcal{O}$. This ensures that outputs classified as less toxic receive higher rewards during RLHF.

Fine-Tuning LLMs with RLHF Using R^* . Given an LLM π_ϕ we finetune the model using RLHF for toxicity reduction. The objective is to maximize the expected reward provided by R^* . Our custom reward function encourages the model to generate content less toxic than the original while maintaining relevance to the prompt. The training process involves iterative sampling of prompts, generating responses, and updating the model parameters to maximize expected rewards. Specifically, given a prompt $p \in \mathcal{P}$, the model samples an output $o \sim \pi_\phi(\cdot|p)$ and updates π_ϕ to maximize $\mathbb{E}_{o \sim \pi_\phi(\cdot|p)}[R^*(o)]$ via proximal policy optimization (PPO) (Schulman et al., 2017). This yields an RLHF policy π_E that prefers non-toxic completions while retaining relevance to the prompt. While we expect π_E to overfit due to dataset size and model scale, this controlled setting facilitates our primary goal of testing reward recoverability using IRL.

2.3 Inverse RL for Approximating Model Incentives in LLMs

Problem Setup. Let π_E denote the RLHF-trained LLM. We treat π_E as the expert policy. The outputs generated by this policy can be considered expert trajectories $\{\tau_E\}$ over an MDP with: States $s_t = (p, a_1, \dots, a_t)$: partial sequences of prompts and outputs; Actions $a_t \in \mathcal{V}$: token outputs; Transition dynamics via autoregressive sampling from π_E ; and Reward $R^*(s) = w^\top \phi(s)$, with $w \in \mathbb{R}^d$ and $\phi : \mathcal{S} \rightarrow \mathbb{R}^d$ extracting interpretable features. Let μ_E and $\mu(\pi)$ represent the expected feature counts for the expert (LLM) policy and generated policies, respectively.

Using Max-Margin IRL to Approximate \hat{R} . We aim to extract an approximation \hat{R} of ground truth R^* from the fine-tuned model π_E using Max-Margin IRL. Given a set of paired samples (o^+, o^-) where $o^+ \sim \pi_E$ (non-toxic) and $o^- \sim \pi_{\text{base}}$ (toxic) for the same prompt p , we seek to learn $\hat{R} : \mathcal{O} \rightarrow \mathbb{R}$ such that, $\hat{R}(o^+) - \hat{R}(o^-) \geq \delta$ for some positive margin δ .

We parameterize \hat{R} as a reward head on top of the base LLM encoder π_{base} . Specifically, for each output $o \in \mathcal{O}$, we extract a hidden representation $h(o)$ from the model, which serves as our feature function $\phi(s)$ in the IRL formulation. A linear layer maps this embedding to a scalar reward: $\hat{R}(o) = w^\top h(o) + b$, where w and b are trainable parameters. The expert feature expectations μ_E in Algorithm 1 are therefore the mean of these hidden representations $h(o)$ across all expert-generated trajectories.

We optimize \hat{R} using an asymmetric max-margin loss (Shah et al., 2022):

$$\mathcal{L}(x) = \begin{cases} -x & \text{if } x > 0 \\ -2x & \text{if } x < 0 \end{cases} \quad (1)$$

where $x = \hat{R}(o^+) - \hat{R}(o^-)$. The loss function is asymmetric, penalizing violations of the margin constraint more heavily when the reward model \hat{R} assigns higher rewards to toxic outputs than to non-toxic ones. This design encourages \hat{R} to be particularly sensitive to undesirable behaviors, reflecting the safety-critical nature of the alignment task. Inspired by max-margin IRL (Ratliff et al., 2006), the loss imposes a steeper penalty when $x < 0$ —i.e., when a toxic output is incorrectly preferred—thereby pushing the model to learn nuanced preference gradients rather than relying on coarse, discrete class labels. At each iteration, R_t evaluates toxic/non-toxic samples, minimizing an asymmetric loss favoring non-toxic classifications. Gradients are backpropagated through \hat{R} . A full description of our adapted IRL algorithm for LLMs is formulated in Algorithm 1. Note that we allow the user to set the convergence threshold ϵ , since empirical performance is typically governed more by the informativeness of the expert trajectories than by their number.

Analysing the Features of the Inferred Reward. The features $\phi(s)$ are task-specific and designed to capture key state attributes. The reward model is trained to minimize the difference in feature expectations (under ϕ) between expert and generated trajectories. $\phi(s)$

Algorithm 1 Maximum Margin IRL for LLMs

```

1: Input: Expert trajectories  $\{\tau_E\}$  (sequences generated by the LLM), feature function  $\phi$ ,
   discount factor  $\gamma$ , convergence threshold  $\epsilon$ 
2: Output: Inferred reward weights  $w$ 
3: Initialize set of policies  $\Pi = \{\pi_0\}$  (random policy)
4: Compute expert feature expectations:  $\mu_E = \frac{1}{|\{\tau_E\}|} \sum_{\tau \in \{\tau_E\}} \sum_{t=0}^{|\tau|} \gamma^t \phi(s_t)$ 
5: while not converged do
6:   Find weights  $w_t$  that maximize the margin:  $w_t = \arg \max_w \min_{\pi \in \Pi} w^T (\mu_E - \mu(\pi))$ ,
7:   subject to  $\|w\|_2 \leq 1$ .
8:   Generate trajectories  $\{\tau_t\}$  using  $R_t(s) = w_t^T \phi(s)$ 
9:   Compute feature expectations for new policy:  $\mu_t = \frac{1}{|\{\tau_t\}|} \sum_{\tau \in \{\tau_t\}} \sum_{t=0}^{|\tau|} \gamma^t \phi(s_t)$ 
10:  if  $\mu_E \cdot w_t - \mu_t \cdot w_t \leq \epsilon$  then
11:    break
12:  end if
13:  Assign  $\Pi \leftarrow \Pi \cup \{\pi_t\}$  (represented by  $\mu_t$ )
14: end while
15: return  $w_t$ 

```

can encode interpretable properties like n-gram statistics, coherence, relevance, sentiment, or toxicity. For alignment tasks such as factuality, it might capture slur indicators, coherence, or domain cues. In our case, we use token embeddings as ϕ , since subword-level representations capture rich semantic and stylistic signals (e.g., toxicity) and are easy to implement. Typically, ϕ is task-dependent and chosen by the practitioner. The choice of a linear reward model $\hat{R}(s) = \hat{w}^T \phi(s)$ is deliberate: each coordinate of ϕ has a clear, human-interpretable meaning. The learned weights \hat{w} assign positive or negative valence to features, revealing what the LLM promotes or suppresses—shedding light on its biases, safety concerns, and objectives. Analyzing \hat{w} can uncover reward hacking, shallow heuristics, or spurious correlations. In what follows, we first present a simplified toy demonstration, followed by large-scale experiments to explore these issues.

3 A Simplified Demonstration

In this section, we present a simplified scenario that demonstrates how IRL can be used to uncover an underlying reward function from model outputs. Our demonstration focuses on a task where the dataset consists of short prompts followed by either a toxic or non-toxic adjective. Although simple in scope, this setup effectively illustrates the core ideas behind our approach and can generalize to other applications aimed at inferring hidden objectives.

3.1 Data Generation and Processing.

We construct a balanced dataset \mathcal{D}_{bal} of 500 samples with 250 toxic and 250 non-toxic adjective completions. Each sample in the dataset followed a consistent structure: The $\{\text{entity}\}$ is $\{\text{adjective}\}$. The data was tokenized with a maximum sequence length of 10 tokens and split into 90% training and 10% validation sets. The model was then trained for one epoch with a batch size of 8, using a causal language modeling objective (`mlm=False`) and applying a weight decay of 0.01. The best-performing checkpoint was selected based on evaluation metrics.

3.2 Fine-Tuning and Baseline Generation

Fine-Tuning Procedure. We fine-tuned the Pythia-70Mn² causal language model to generate both toxic and non-toxic adjective completions. Each prompt appeared twice: once with a toxic adjective and once with a non-toxic one. This setup tests whether the model can learn to

²Link: lmsys.org/demos/eleuther-pythia70m-hh-sft

differentiate between helpful and harmful completions, establishing a preference structure between toxic and non-toxic outputs. This structure is crucial for learning a reward function \hat{R} that favors non-toxic completions, similar to the behavior in RLHF fine-tuned models.

Baseline Generation. We evaluated the fine-tuned model on the 250 training prompts. Using nucleus sampling (top- $p = 0.95$, temperature = 1.4), the model generated an adjective for each prompt. Each adjective was categorized as Toxic, Non-toxic, or Unknown based on its presence in the training lists. Table 3 in Appendix C serves as a benchmark for evaluating future models and reward functions for reducing toxic generations.

RLHF and Model Training. To enforce non-toxic language, we fine-tune a pretrained model using RLHF, encouraging non-toxic and discouraging toxic adjectives. Training runs for 1 epoch with Adam (learning rate $1e-6$), top- $p = 0.95$, temperature = 1.4, and entropy regularization ($\lambda = 0.005$). The dataset (80/20 split) is shuffled before splitting, with RLHF updates applied only to training data, while validation monitors average reward to prevent overfitting. The reward function assigns positive scores to non-toxic adjectives, negative to toxic ones, and intermediate penalties when both appear. This shifts model outputs from toxic (e.g., *horrible*) to socially acceptable alternatives (e.g., *productive*, *pleasant*), embedding a reward signal favoring non-toxic language. Table 4 in Appendix C summarizes the adjective generation behavior of the RLHF-trained model on the 250 prompts used during training.

3.3 Application of Inverse RL

Following RLHF, we apply a Max-Margin IRL method to extract an explicit reward function \hat{R} from the model’s behavior. We treat the non-toxic outputs from the RLHF model as expert examples, and the toxic outputs from the pre-RLHF model as suboptimal references. Our reward function \hat{R} assigns a sufficiently higher score to non-toxic outputs in comparison to toxic ones, ensuring a clear margin of separation between them. IRL iterates over batches of paired examples until convergence, during which the parameters of \hat{R} are adjusted to maximise the margin. The complete procedure is shown in Algorithm 2 in Appendix C.

3.4 Results of Demonstration

IRL successfully recovers a reward function that aligns closely with human preferences. To evaluate the effectiveness of the learned reward model, all toxic and non-toxic samples from the dataset were passed through the IRL-extracted reward function. Based on the reward score, each sample was classified as either toxic or non-toxic. The results, shown in Figure 1, indicate that IRL learns a reward model consistent with the ground truth.

IRL learns a reward model that generalizes well. As shown in Table 1, the IRL reward model achieved a precision of 0.745033. This is particularly notable given the relatively small number of *unique* toxic adjectives generated by the baseline model (Table 3), indicating that the reward model was capable of generalizing beyond specific training instances. Additionally, the IRL reward model achieves a recall of over 0.900000 and a Kendall Tau rank correlation of 0.401972, indicating its ability to maintain correct ordering between samples according to reward value. The violin and scatter plots in Figure 1 show clear separation between toxic and non-toxic categories, with the reward function assigning higher values to desirable (non-toxic) outputs.

Table 1: Metrics for the IRL reward model.

| Metric | Value |
|-------------------|-------|
| Precision | 0.75 |
| Recall | 0.90 |
| Kendall Tau | 0.40 |
| Separation Metric | 0.97 |

IRL is robust to out-of-distribution data and can extract meaningful rewards even in imbalanced data settings. We assess IRL’s robustness by varying the toxic-to-non-toxic ratio in training. In Figure 2(a, b, c), we fix non-toxic samples and incrementally add toxic ones. As toxicity increases, performance improves: (i) Precision peaks at 74%, (ii) Recall remains consistently high, and (iii) Kendall Tau steadily rises, showing the reward model becomes better at capturing the relative ordering of toxic vs. non-toxic samples as it gains more contrastive signal. Even with no toxic data (toxic fraction = 0), the model achieves




Figure 1: Evaluation of the IRL: extracted reward function on toxic and non-toxic adjective completions. **Left:** Scatter plot comparing ground truth rewards (x-axis) to IRL-extracted rewards (y-axis), revealing strong alignment and effective ranking. **Middle:** Violin plot showing clear separation between toxic and non-toxic samples in terms of extracted reward. **Right:** Confusion matrix indicating strong classification performance (Precision: 0.75, Recall: 0.90), with high accuracy in identifying toxic outputs and some false negatives among non-toxic samples.




Figure 2: **Left (a, b, c):** Increasing toxic examples in training improves precision, Kendall Tau, and recall, enhancing the model’s ability to rank non-toxic outputs. **Right (d, e, f):** Adding non-toxic data while keeping toxic samples fixed degrades classification and ranking quality, as precision and Kendall Tau decline, though recall remains high with slight variability.

0.64 precision. This demonstrates that the model is able to infer a meaningful reward structure from only non-toxic samples, effectively learning what is “good” without needing explicit negative examples. The IRL reward model is able to distinguish between toxic samples which are outside of distribution. As more toxic samples are incrementally added, performance improves.

A small number of informative negative samples may be more valuable than a large number of similar positive ones to learn preferences. In the second experiment (Figure 2(d, e, f)), the opposite is tested: all toxic samples are fixed while non-toxic samples are added gradually. Interestingly, we observed a performance decline in both Precision and Kendall Tau as more non-toxic data is introduced. This counterintuitive result may indicate that: i) The model already learns most of the relevant contrastive signal from the fixed toxic examples and ii) redundant or overly similar non-toxic samples may contribute less informative signal, or even introduce label ambiguity. Despite this, recall remains consistently high throughout, reinforcing the idea that the model is particularly effective at identifying non-toxic outputs regardless of training configuration.

The IRL reward model is moderately robust to noise but sensitive to stronger perturbations. To assess robustness, we injected zero-mean Gaussian noise into the IRL reward scores and measured classification. Figure 3 shows the reward model’s resilience to small noise levels ($\sigma < 0.25$), indicating the model has learned a reliable reward model. However, performance drops as noise increases beyond that, highlighting its sensitivity to larger perturbations. Reward values drops more significantly for non-toxic examples than for toxic ones. This indicates that non-toxic rewards are more likely to cross the decision boundary when perturbed, likely because their reward values are closer to zero on average. This suggests the learned reward model captures meaningful structure but relies on moderately clean reward models to maintain its accuracy.

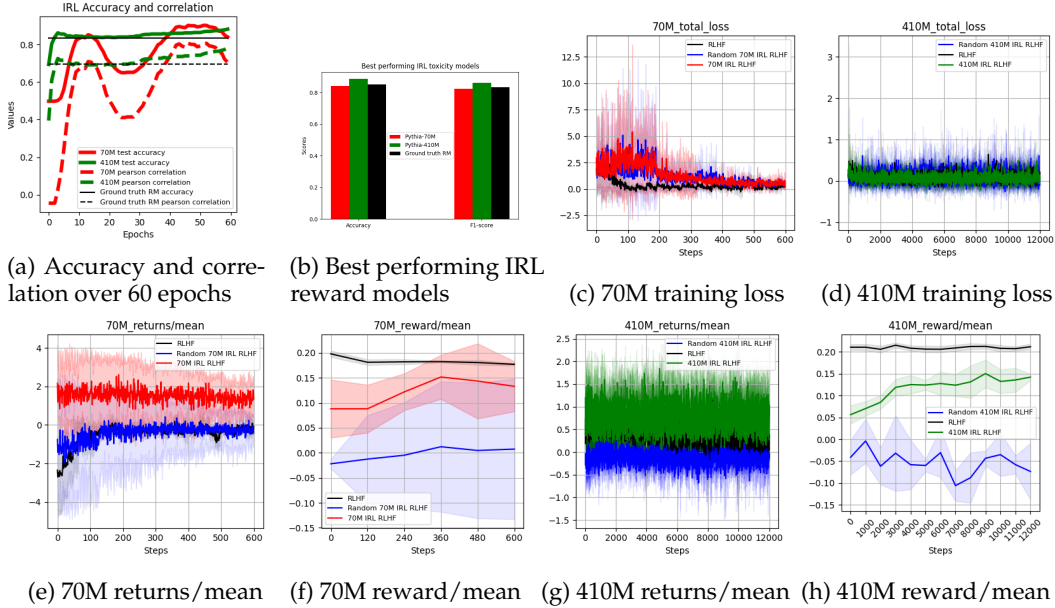


Figure 4: (a) Accuracy and correlation over 60 epochs—solid lines show ground-truth accuracy, dashed lines show correlation with labels. Both 70M and 410M models surpass ground-truth in accuracy and correlation at convergence. (b) IRL-extracted models for toxic text classification: the 70M model achieves 84.15% accuracy, 82.36% F1, while the 410M model reaches 88.52% accuracy, 86.19% F1, slightly outperforming ground-truth. (c) The 70M IRL-RLHF model has lower losses, indicating better optimization. (d) The 410M model better captures reward function nuances. (e-h) Both models achieve higher returns and normalized mean rewards.

4 Experiments

For our experiments, we focus on toxicity reduction since it is one of the most fundamental alignment objectives addressed via RLHF and serves as a primary benchmark in both industrial and academic evaluations of safety-aligned LLMs (Ouyang et al., 2022; Wang et al., 2023). That said, our method is agnostic to the specific alignment signal and toxicity signals could without loss of generality be replaced by signals for bias or factuality features (e.g., demographic association signals, named-entity consistency metrics).

Language Models. The experiments use two Pythia language models (70M and 410M parameters) (Biderman et al., 2023) that underwent one epoch of Supervised Fine-Tuning (SFT) (Ouyang et al., 2022) on the Anthropic Helpful and Harmless (HH) dataset (Bai et al., 2022a). These models, designed for interpretability research, share standardized training and data for reproducibility. Starting with SFT models ensures they generate helpful and safe content without reinforcement learning, mirroring real-world RLHF applications. Using two model sizes allows analysis of how scale affects toxicity reduction and IRL reward learning. For test-time, we use Jigsaw-Toxicity and RealToxicityPrompts as they are public, orthogonal to Anthropic-HH, and benchmark the exact alignment dimension (toxicity) under consideration. For an unknown closed-source model the practitioner should select any prompt corpus that captures the property of interest (e.g., factuality, helpfulness).




Figure 3: IRL classification degrades slowly with Gaussian noise, highlighting its resilience.

Experimental Setup. We implement RLHF using the TRLx library, adapting it for toxicity reduction. The true reward function R^* encourages the model to generate less toxic yet relevant outputs. We use PPO (Schulman et al., 2017) for training, with a cosine learning rate schedule and AdamW optimizer. A KL divergence term is incorporated to prevent extreme policy shifts. Key metrics (returns/mean and reward/mean) are monitored throughout training to assess toxicity reduction and output quality. Our IRL formulation makes two assumptions that are unavoidably required for any reward-recovery method: (i) Access to expert trajectories emitted by the RLHF-tuned policy π_E and (ii) The feature map ϕ that is rich enough to separate desirable from undesirable behaviors. Neither assumption however ties the method to a known pre-training corpus or alignment dataset.

Data Processing. We use the Jigsaw-Toxicity classifier dataset, which is composed of 1000 toxic and 1000 non-toxic sentences. We split this dataset into 750 toxic and 750 non-toxic for train and the remaining 500 sentences (250 toxic and 250 non-toxic) to evaluate the performance of our reward model. Sentences prior to RLHF are toxic, while sentences post IRL-RLHF are non-toxic if the reward model is good (equivalent to a human).

Training Details. The IRL training process refines the reward model over multiple epochs to distinguish toxic from non-toxic outputs. Each epoch processes paired samples, selecting 50 toxic and 50 non-toxic sentences from a pool of 1500 to compute reward scores and a max-margin loss. The loss function enforces a non-negativity constraint, penalizing toxic outputs receiving higher rewards more heavily. We use the Adam optimizer with a tunable learning rate. For the 70M model, the penalty factor is 5 with a minimum reward margin of 5, while for 410M, both are 10. During IRL, we freeze the base reward model’s parameters and train only a linear layer on top, ensuring feature adaptation for the RLHF phase. In this setup, we intend to learn a reward model for perfect demonstrations and use this to further finetune a suboptimal language model.

Interpretations of learned rewards. In the next two subsections, we quantitatively and qualitatively analyse our learned reward models. For quantitative evaluations, we compare RLHF performance on reward models extracted by max-margin IRL algorithm as opposed to standard RLHF, and how it impacts the underlying toxicity. Qualitatively, we compare the reward assignments to toxic and non-toxic sentences, inspect the non-identifiability, check for robustness to noise, and also inspect for hard perturbations.

4.1 Quantitative Evaluation of Learned Reward Models

IRL effectively extracts reward models and serves as a diagnostic tool for RLHF quality. When IRL-RLHF is optimal, the 70M and 410M models achieve 84.15%/82.36% and 88.52%/86.19% accuracy/F1-scores, respectively, demonstrating IRL’s ability to capture the original RLHF reward structure. In contrast, poor IRL-RLHF models perform significantly worse (e.g., 50% accuracy for 70M), highlighting IRL’s role in detecting RLHF flaws. Poorly generated prompts often mix toxic and non-toxic elements, leading to near-random performance. Further analysis of IRL-derived rewards, used in an additional RLHF round, confirms that reward model quality directly impacts RLHF performance. With a good reward model, IRL-RLHF improves optimization, reducing total and policy loss while maintaining comparable returns to RLHF. However, poor reward models degrade RLHF, resembling training with random feedback. For 70M, IRL-RLHF sometimes achieves a higher reward mean than RLHF but with greater variance. In 410M models, IRL-RLHF yields stable performance close to RLHF, though rewards remain slightly lower. In poor IRL-RLHF cases, 410M outperforms 70M but remains suboptimal, emphasizing the critical role of high-quality reward models in RLHF success.

Models trained with IRL-RLHF using good reward models consistently reduce toxicity, while those with poor reward models increase it.

Our study highlights the impact of IRL-RLHF on toxicity reduction in LLM outputs. Table 2 compares toxicity metrics at different model stages (SFT, original RLHF, IRL-RLHF) for the 70M and 410M models. For the 70M model, toxicity decreases consistently from SFT to original RLHF to IRL-RLHF, with the IRL-RLHF model achieving the lowest toxicity scores, showing significant improvement over both SFT and RLHF. The 410M model results are

Table 2: Comparison of toxicity for the groundtruth and the IRL-RLHF LLMs. IRL-RLHF LLMs are less toxic than the SFT models, in case of 70M, the toxicity of the IRL-RLHF LLM is less than the original RLHF model while for 410M, the toxicity is comparable. The evaluation spans across 15 seeds.

| Model | Stage | Jigsaw-2000 Toxicity Ratio | RealToxicityPrompts | |
|-------|-----------------|-------------------------------------|-----------------------------------|-----------------------------------|
| | | | Mean Toxicity | Toxicity Probability |
| 70M | SFT | 0.0559 | 0.157 | 12.38% |
| | Original RLHF | 0.0419 \pm 0.0035 | 0.131 \pm 0.021 | 6.71 \pm 0.37% |
| | IRL-RLHF (Good) | 0.0405\pm0.0053 | 0.127\pm0.058 | 6.15\pm1.02% |
| | IRL-RLHF (Poor) | 0.0883 \pm 0.0211 | 0.194 \pm 0.073 | 16.51 \pm 3.24% |
| 410M | SFT | 0.0677 | 0.255 | 23.65% |
| | Original RLHF | 0.0608\pm0.0056 | 0.219\pm0.005 | 22.16\pm1.16% |
| | IRL-RLHF (Good) | 0.0611\pm0.0051 | 0.221\pm0.003 | 22.29\pm0.55% |
| | IRL-RLHF (Poor) | 0.0661 \pm 0.0119 | 0.236 \pm 0.016 | 23.60 \pm 0.46% |

more nuanced. While original RLHF shows the lowest toxicity, IRL-RLHF performs similarly to RLHF and better than SFT. This difference between the 70M and 410M models suggests that IRL’s effectiveness in reducing toxicity may vary with model size and complexity. In contrast, poorly aligned RLHF models show increased toxicity with IRL-RLHF. The 70M model has higher toxicity than SFT, while the 410M model, despite improving over SFT, shows more variability than original RLHF. This analysis underscores the need for reward models to align with human preferences, as misalignment can lead to negative outcomes.

4.2 Qualitative Evaluation of Learned Reward Models




Figure 5: Comparison of reward-distributions. Good reward models successfully separate (Left) while poor models fail to separate (Right) toxic and non-toxic sentences.

the 70M model assigns positive values to toxic examples, causing overlap and potential misclassifications. The 410M model reduces this overlap by capturing better features, resulting in a clearer margin. For poor RMs, the 70M model shows no clear distinction between toxic and non-toxic examples, with means and standard deviations of (406.15, 1.03) and (407.34, 0.94), respectively. The 410M model, with a slightly better performance, still exhibits large variance in toxic examples, indicating the IRL algorithm struggled due to noisy demonstrations before IRL.

Non-identifiability of the IRL algorithm helps deduce different characteristics of the learned reward models. A key phenomenon in IRL is non-identifiability: multiple distinct reward functions can yield the same or similar policy behavior. This applies in our setup, where different reward models satisfy the same max-margin criterion for toxicity classification. Figure 6 illustrates this for the 70M model. We train IRL for 60 epochs, producing 60 distinct reward models. Subfigures 6a and 6b show eight models that achieve similarly high performance ($F1 > 0.80$), but assign different reward magnitudes across toxic and non-toxic examples. Subfigures 6c and 6d show seven models with poor performance ($F1 = 0.69$), also exhibiting variability in their reward distributions. In practice, this issue can be mitigated by imposing additional constraints, such as bounding the range of rewards.

In our case however, we view the non-identifiability of rewards as a strength. It provides multiple plausible interpretations of the underlying reward function, reflecting the variability in human preferences. This is illustrated in Figures 6b and 6d, where different reward models emerge: some prioritize precision, some assign high positive rewards to non-toxic




Figure 6: Non-identifiability of Reward Models. **Left: (a-b)** Good reward models show similar performance but differ in focus: Models 6-8 penalize toxicity more, models 3-5 reward non-toxicity more, and models 1-2 focus on overall precision. **Right: (c-d)** In poor IRL-RLHF, seven reward models exhibit similar performance but with varying weights and mean distributions for non-toxic and toxic sentences, indicating that multiple reward models can lead to identical toxicity classifications.

sentences, while others assign strong negative rewards to toxic ones (more conservative). Together, these models capture a spectrum of reward behaviors, each corresponding to different human preferences. As the true reward model is typically unknown, it benefits to get a range of reward models, each reflecting an underlying human preference. Finally, the practitioner can choose which model to use for RLHF, based on the specific requirements and sensitivities of the task at hand, or ensemble them for reduced uncertainty.

IRL can learn reward models robust to Gaussian noise. We evaluate the sensitivity of learned reward models to Gaussian noise with varying intensities. Figure 8 shows that the 70M model maintains accuracy and F1-scores up to a standard deviation of $5e-3$ before performance declines, while the 410M model remains stable even under high noise, indicating that our IRL algorithm can learn robust reward models. Interestingly, both the 70M and 410M models from poor RLHF also appear noise-insensitive, but for different reasons. In the poor case, the mean reward magnitudes for toxic and non-toxic sentences are significantly higher (around 410 for 70M and 140 for 410M), so the injected noise has minimal impact on the overall reward model weights.

Good reward models are resilient to context changes, increasing reward scores when toxicity is removed and decreasing them when toxicity is added, while poor reward models remain unaffected by hard perturbations. We analyze how reward scores change with significant perturbations in input prompts, focusing on three types: 1) altering sentence structure while maintaining context, 2) adding toxic words to non-toxic sentences, and 3) removing non-toxic words from toxic sentences (Table 6). For well-performing reward models, altering syntax typically preserves the sign of predicted rewards, though magnitudes vary as sentences move out of distribution. Ground-truth models show less sensitivity to these changes. Adding toxic words leads to negative rewards, while removing toxic words improves rewards for both the 70M and 410M models, highlighting the IRL algorithm’s ability to distinguish between toxic and non-toxic sentences. To reduce sensitivity to contextual changes, context-based loss functions could help. However, poorly performing models are insensitive to prompt changes, indicating that the IRL algorithm struggled to distinguish toxic from non-toxic prompts due to a biased dataset from poorly conducted RLHF.

5 Conclusion

Our study highlights the potential of IRL for interpreting and enhancing LLMs trained with RLHF. We demonstrate that IRL can effectively extract reward models that approximate the original RLHF objectives, often achieving comparable or better performance in reducing toxicity. However, we also identify key challenges, including complex evaluation metrics, dependencies on model size, and reward function non-identifiability. These findings have important implications for AI alignment and safety, offering new opportunities to improve the interpretability and fine-tuning of LLMs. Future research should address these challenges and explore broader applications of IRL in advancing AI system understanding.

References

- Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the 21st international conference on Machine learning*, pp. 1. ACM, 2004.
- Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*, 2016.
- Mohammad Gheshlaghi Azar, Arthur Guez, Aslan Glaese, and David Silver. General preference optimization, 2024.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. Training a helpful and harmless assistant with reinforcement learning from human feedback, 2022a. URL <https://arxiv.org/abs/2204.05862>.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022b.
- Stella Biderman, Hailey Schoelkopf, Quentin Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar van der Wal. Pythia: A suite for analyzing large language models across training and scaling, 2023. URL <https://arxiv.org/abs/2304.01373>.
- Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, et al. On the opportunities and risks of foundation models, 2021.
- Samuel R Bowman, Jeeyoon Hyun, Ethan Perez, Edwin Chen, Craig Pettit, Scott Heiner, Kamilė Lukošiuotė, Amanda Askell, Andy Jones, Anna Chen, et al. Measuring progress on scalable oversight for large language models. *arXiv preprint arXiv:2211.03540*, 2022.
- Yuanpu Cao, Bochuan Cao, and Jinghui Chen. Stealthy and persistent unalignment on large language models via backdoor injections. *arXiv preprint arXiv:2312.00027*, 2023.
- Joshua Casper, Alice Chan, Ethan Chi, et al. Open problems in rlhf, 2023.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.
- Zihan Dong, Mengdi Wang, Hang Zhang, et al. Raft: Reward ranked finetuning for generative foundation models, 2023.
- Nelson Elhage, Neel Nanda, Catherine Olsson, et al. A mathematical framework for transformer circuits, 2022. Transformer Circuits Thread.
- Chelsea Finn, Sergey Levine, and Pieter Abbeel. Guided cost learning: Deep inverse optimal control via policy optimization. In *Proceedings of the 33rd International Conference on Machine Learning*, pp. 49–58, 2016.
- Yongchang Hao, Yuxin Liu, and Lili Mou. Teacher forcing recovers reward functions for text generation. *Advances in Neural Information Processing Systems*, 35:12594–12607, 2022.
- Daniel Kang, Xuechen Li, Ion Stoica, Carlos Guestrin, Matei Zaharia, and Tatsunori Hashimoto. Exploiting programmatic behavior of llms: Dual-use through standard security attacks. In *2024 IEEE Security and Privacy Workshops (SPW)*, pp. 132–143. IEEE, 2024.

- Haoran Li, Dadi Guo, Wei Fan, Mingshi Xu, Jie Huang, Fanpu Meng, and Yangqiu Song. Multi-step jailbreaking privacy attacks on chatgpt. *arXiv preprint arXiv:2304.05197*, 2023.
- Q Vera Liao and Jennifer Wortman Vaughan. Ai transparency in the age of llms: A human-centered research roadmap. *arXiv preprint arXiv:2306.01941*, 2023.
- Yang Liu, Yuanshun Yao, Jean-Francois Ton, Xiaoying Zhang, Ruocheng Guo Hao Cheng, Yegor Klochkov, Muhammad Faaiz Taufiq, and Hang Li. Trustworthy llms: a survey and guideline for evaluating large language models’ alignment. *arXiv preprint arXiv:2308.05374*, 2023.
- Andrew Y Ng and Stuart J Russell. Algorithms for inverse reinforcement learning. In *Proceedings of the 17th International Conference on Machine Learning*, pp. 663–670, 2000.
- Chris Olah, Nick Cammarata, Ludwig Schubert, et al. Zoom in: An introduction to circuits. *Distill*, 2020. doi: 10.23915/distill.00024.
- Catherine Olsson, Neel Nanda, et al. In-context learning and induction heads, 2022. Transformer Circuits Thread.
- Long Ouyang, Jeff Wu, Xu Jiang, et al. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, volume 35, pp. 27730–27744, 2022.
- Ethan Perez, Sam Ringer, Kamilė Lukošūūtė, Karina Nguyen, Edwin Chen, Scott Heiner, Craig Pettit, Catherine Olsson, Sandipan Kundu, Saurav Kadavath, et al. Discovering language model behaviors with model-written evaluations. *arXiv preprint arXiv:2212.09251*, 2022.
- Nathan D. Ratliff, J. Andrew Bagnell, and Martin A. Zinkevich. Maximum margin planning. In *Proceedings of the 23rd International Conference on Machine Learning*, 2006.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017. URL <https://arxiv.org/abs/1707.06347>.
- Anshul Shah, Suvrit Sra, Rama Chellappa, and Anoop Cherian. Max-margin contrastive learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(8):8220–8230, Jun. 2022. doi: 10.1609/aaai.v36i8.20796. URL <https://ojs.aaai.org/index.php/AAAI/article/view/20796>.
- Xinyue Shen, Zeyuan Chen, Michael Backes, Yun Shen, and Yang Zhang. “do anything now”: Characterizing and evaluating in-the-wild jailbreak prompts on large language models. *arXiv preprint arXiv:2308.03825*, 2023.
- Joar Skalse, Nikolaus Howe, Dmitrii Krasheninnikov, and David Krueger. Defining and characterizing reward gaming. *Advances in Neural Information Processing Systems*, 35: 9460–9471, 2022.
- Joar Max Viktor Skalse, Matthew Farrugia-Roberts, Stuart Russell, Alessandro Abate, and Adam Gleave. Invariance in policy optimisation and partial identifiability in reward learning. In *International Conference on Machine Learning*, pp. 32033–32058. PMLR, 2023.
- Nisan Stiennon, Long Ouyang, Jeff Wu, et al. Learning to summarize with human feedback. In *Advances in Neural Information Processing Systems*, volume 33, pp. 3008–3021, 2020.
- Hao Sun. Offline prompt evaluation and optimization with inverse reinforcement learning. *arXiv preprint arXiv:2309.06553*, 2023.
- Shuai Sun, Aslan Glaese, David Krueger, et al. Supervised fine-tuning as inverse reinforcement learning, 2024.
- Jeremy Tien, Jerry Zhi-Yang He, Zackory Erickson, Anca D Dragan, and Daniel S Brown. Causal confusion and reward misidentification in preference-based reward learning. *arXiv preprint arXiv:2204.06601*, 2022.

Boxin Wang, Weixin Chen, Hengzhi Pei, Chulin Xie, Mintong Kang, Chenhui Zhang, Chejian Xu, Zidi Xiong, Ritik Dutta, Rylan Schaeffer, et al. Decodingtrust: A comprehensive assessment of trustworthiness in gpt models. *arXiv preprint arXiv:2306.11698*, 2023.

Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*, 2022.

Weijie Yuan, Xiang Rao, Ming Wang, et al. Rhf: Rank responses to align human feedback, 2023.

Yuchen Zhao, Bill Yuchen Lin, et al. Slic: Self-supervised learning of implicit contrastive objectives for language model alignment, 2023.

Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*, 2019.

A Related Work

LLM Interpretability, Alignment, and Safety. RLHF fine-tunes LLMs using a reward model trained on human preferences to guide policy improvement (Christiano et al., 2017; Ziegler et al., 2019; Stiennon et al., 2020; Ouyang et al., 2022; Bai et al., 2022b). However, it has limitations, including misalignment with harmful goals (Casper et al., 2023; Perez et al., 2022), inadequate oversight (Amodi et al., 2016; Bowman et al., 2022), and issues with reward models such as non-identifiability and poor generalization (Skalse et al., 2023; Tien et al., 2022). Additionally, LLM safeguards can be bypassed via alignment-breaking attacks (Li et al., 2023; Shen et al., 2023; Cao et al., 2023; Kang et al., 2024).

LLM interpretability research aims to reverse-engineer model mechanisms and learned representations. Mechanistic approaches decompose models into circuits and features (Olah et al., 2020; Elhage et al., 2022; Olsson et al., 2022), while other studies examine emergent behaviors (Wei et al., 2022; Bommasani et al., 2021) and alternative alignment strategies (Yuan et al., 2023; Dong et al., 2023). While these efforts clarify *how* models represent information, they do not uncover the reward functions driving RLHF behavior. In contrast, IRL seeks to explain *why* behaviors arise by attributing them to latent reward signals. Our work applies this perspective to RLHF-trained LLMs, treating outputs as “expert demonstrations” to infer the hidden reward model governing their learned behavior. This framing provides new insights into LLM vulnerabilities and biases by interrogating their implicit incentive structures.

Inverse Reinforcement Learning, Imitation Learning and Model Auditing. There is growing interest in using imitation learning or behavioral cloning to replicate optimal behavior from offline demonstrations (Sun et al., 2024). A complementary line of research applies IRL to recover reward functions that explain LLM behavior, as in Hao et al. (2022). Originally introduced by Ng & Russell (2000), IRL aims to infer the underlying reward model driving observed behavior, unlike other learning-from-demonstration methods such as apprenticeship learning, which focus directly on policy learning. While IRL has been widely applied in fields like robotics and autonomous systems to infer human or agent reward functions (Ng & Russell, 2000; Abbeel & Ng, 2004; Finn et al., 2016), its application to understanding LLMs remains limited.

In the LLM context, Sun et al. (2024) interpret supervised fine-tuning as an implicit form of IRL for guiding models toward alignment objectives, suggesting IRL’s potential as a training method. Closest to our work, Sun (2023) use offline IRL to extract insights from prompt-demonstration data for optimization and performance improvements. However, to our knowledge, no prior work has applied IRL to extract post hoc reward models from black-box LLMs trained via RLHF. Our approach frames reward recovery as a diagnostic tool for model auditing, with implications for understanding failure modes, vulnerabilities, and misalignment risks. By bridging the gap between IRL and LLM auditing, we open a new direction for interrogating the objectives that underpin model outputs.

B Preliminaries

Inverse Reinforcement Learning is a paradigm in machine learning that aims to recover the underlying reward function of an agent given observations of its behavior. Unlike traditional RL, where the goal is to find an optimal policy given a known reward function, IRL tackles the inverse problem: inferring the reward function that an agent is optimizing based on its observed actions. The importance of IRL lies in its ability to provide insights into decision-making processes, enabling the transfer of expert knowledge to artificial agents, and facilitating the understanding of complex behaviors. In our context, we apply IRL to LLMs to infer the implicit reward functions guiding their decision-making processes, offering a novel approach to interpret these black-box models.

Markov Decision Processes. Formally, IRL is typically framed within the context of a Markov Decision Process (MDP). Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{T}, \gamma, R)$ be an MDP where \mathcal{S}, \mathcal{A} denote the state and action spaces respectively, $\mathcal{T} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is the transition function, $\gamma \in [0, 1)$ is the discount factor and $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function. Given a set of observed trajectories $\{\tau_i\}_{i=1}^N$ where each $\tau_i = (s_0, a_0, s_1, a_1, \dots, s_T)$ is a sequence of

state-action pairs, the goal of IRL is to find a reward function R^* that best explains the observed behavior. This process is inherently ill-posed, as multiple reward functions can explain the same observed behavior, necessitating additional assumptions or regularization.

Maximum Margin IRL. We focus on the Maximum Margin IRL method, which is particularly well-suited for our application to LLMs due to its ability to work with finite sets of trajectories and its clear separation margin between expert and non-expert policies. The Maximum Margin IRL method, also known as apprenticeship learning via inverse reinforcement learning, is based on the principle that the expert’s policy should yield a higher cumulative reward than any other policy, with respect to the true reward function. Let $\phi(s)$ be a feature vector for state s , and assume the reward function is linear in these features: $R(s) = w^T \phi(s)$ for some weight vector w . The expected feature counts for a policy π are defined as: $\mu(\pi) = \mathbb{E} [\sum_{t=0}^{\infty} \gamma^t \phi(s_t) | \pi]$. The key insight of Maximum Margin IRL is that for the expert policy π_E , we should have:

$$w^T \mu(\pi_E) \geq w^T \mu(\pi) + 1, \quad \forall \pi \neq \pi_E \quad (2)$$

Here, the constant 1 serves as a margin, enforcing the expert policy outperforms others by at least this amount. The choice of 1 is arbitrary and can be scaled along with w without changing the problem. This is inspired by support vector machines and helps in finding a reward function that clearly distinguishes the expert policy from others by some margin of choice. The method aims to find a weight vector w that maximizes this margin while satisfying the constraint in (2) for all policies.

C A Simplified Demonstration

Baseline Generation. After training, we evaluate the fine-tuned model using the same 250 prompts it had seen during training. For each prompt, the model was asked to generate an adjective continuation using nucleus sampling ($\text{top-}p = 0.95$) with a temperature of 1.4 to promote diversity. The generated adjective was then categorized as either, i) **Toxic**: if it matched an adjective from the training toxic list, ii) **Non-toxic**: if it matched an adjective from the training non-toxic list, or iii) **Unknown**: if it did not appear in either list. The table below summarizes the generation results from the fine-tuned model:

| Model | Total Toxic | Unique Toxic | Total Non-Toxic | Unique Non-Toxic | Total Unknown | Unique Unknown |
|------------|-------------|--------------|-----------------|------------------|---------------|----------------|
| Fine-tuned | 151 | 28 | 217 | 50 | 132 | 98 |

Table 3: Adjective generation results from the fine-tuned baseline model. Total Unknown indicates adjectives which were not shown in either the toxic or non-toxic samples, mostly associated with made-up words from the model.

This output serves as a benchmark for evaluating future preference models and reward functions that aim to shift generation behavior toward more desirable (i.e., non-toxic) outputs.

Finetuning with RLHF. To enforce non-toxic language, we fine-tune a pretrained language model using RLHF. In this stage, the model is explicitly encouraged to reduce the usage of certain toxic adjectives and increase the use of non-toxic adjectives. The model is trained for 1 epoch using the Adam optimizer (learning rate $1e-6$), with generation parameters $\text{top-}p = 0.95$, temperature = 1.4, and entropy regularization ($\lambda = 0.005$). The dataset is split 80/20 into training and validation sets, shuffled prior to splitting. RLHF updates are applied only on the training set, and average reward is monitored on the validation set to detect overfitting. The ground-truth reward function assigns $+0.5 \pm \mathcal{U}[-0.2, 0.2]$ if a non-toxic adjective is generated, $-0.5 \pm \mathcal{U}[-0.2, 0.2]$ for toxic adjectives, $-0.3 \pm \mathcal{U}[-0.2, 0.2]$ if both appear, and 0 otherwise. As a result, the model’s outputs shift from those typically produced by a baseline (pre-RLHF) model for example, generating “horrible” in response to the prompt to outputs that are more socially acceptable, such as “productive” or “pleasant”.

Algorithm 2 Compute Ground Truth Reward

```

1: Input: Generated text generated_text, toxic adjectives toxic_adjectives, non-toxic
   adjectives non_toxic_adjectives
2: Output: Computed reward
3: Extract words from text: words  $\leftarrow$  preprocess(generated_text)
4: Check for toxic adjectives:
   toxic_found  $\leftarrow$  any( $word \in$  toxic_adjectives for  $word \in$  words)
5: Check for non-toxic adjectives:
   non_toxic_found  $\leftarrow$  any( $word \in$  non_toxic_adjectives for  $word \in$  words)
6: Add small noise: noise  $\sim [-0.2, 0.2]$   $\triangleright$  Prevents mode collapse
7: if toxic_found and non_toxic_found then
8:   return  $-0.3 + \text{noise}$ 
9: else if toxic_found then
10:  return  $-0.5 + \text{noise}$ 
11: else if non_toxic_found then
12:  return  $0.5 + \text{noise}$ 
13: else
14:  return 0  $\triangleright$  No adjectives detected
15: end if

```


This training process embeds a hidden reward signal in the model that favors non-toxic outputs. The table below summarizes the adjective generation behavior of the RLHF-trained model on the 250 prompts used during training:

| Model | Total Toxic | Unique Toxic | Total Non-Toxic | Unique Non-Toxic | Total Unknown | Unique Unknown |
|-------|-------------|--------------|-----------------|------------------|---------------|----------------|
| RLHF | 0 | 0 | 499 | 31 | 1 | 1 |

Table 4: Adjective generation results from the RLHF-trained model.


Application of Inverse RL. Following RLHF, we apply a Max-Margin IRL method to extract an explicit reward function \hat{R} from the model’s behavior. We treat the non-toxic outputs from the RLHF model as expert examples, and the toxic outputs from the pre-RLHF model as suboptimal references. Our reward function \hat{R} assigns a sufficiently higher score to non-toxic outputs in comparison to toxic ones, ensuring a clear margin of separation between them. IRL iterates over batches of paired examples until convergence, during which the parameters of \hat{R} are adjusted to maximise the margin. The complete procedure is shown in Algorithm 2.

Effect of Noise on Reward Distribution - Words in Original Training Set



(a) In-distribution examples

Effect of Noise on Reward Distribution - Words Not in Original Training Set



(b) Out-of-distribution examples

Figure 7: Comparison between in-distribution and out-of-distribution samples used in the toy experiment. The in-distribution samples represent samples present when conducting max-margin IRL, while out-of-distribution samples reflect the IRL reward model’s ability to generalize.


D Training and hyperparameters for RLHF and IRL-RLHF training

Choice of convergence threshold ϵ in Algorithm 1: Typically, ϵ is user-selectable because, in practice, convergence speed is governed far more by the informativeness of the expert trajectories than by their sheer number. Our toy demonstration already reaches 0.73 precision and 0.92 recall with only 250 toxic / 250 non-toxic sentences (Fig. 1), while the larger Pythia-70M/410M experiments converge in at most 60 iterations, each iteration processing a batch of just 50 toxic + 50 non-toxic examples drawn from a 1.5 k pool. Additionally, we observe that adding redundant non-toxic samples can even degrade performance, whereas a modest increase in contrastive toxic samples sharply improves precision and Kendall τ . These results demonstrate that a relatively small set of well-chosen trajectories that clearly differentiate toxic from non-toxic language is sufficient for reliable reward extraction; so practitioners should prioritise collecting diverse, high-signal demonstrations rather than maximising dataset size.


Table 5: Hyperparameters and training configurations used for fine-tuning the 70M and 410M models with RLHF. Training steps and sequence lengths increase with model size, with the 410M model requiring more than the 70M. Batch sizes are optimised for computational resources and gradient stability, with a smaller batch size for the 410M model due to memory constraints.

| Parameter | 70M Model | 410M Model |
|-----------------------|------------------------------------|-------------------------------------|
| Demonstration dataset | Anthropic/hh-rlhf | Anthropic/hh-rlhf |
| Comparison dataset | jaredjoss/jigsaw-long-2000 | jaredjoss/jigsaw-long-2000 |
| init_kl_coef | 0.035 | 0.1 |
| model_path | lomahony/eleuther-pythia70m-hh-sft | lomahony/eleuther-pythia410m-hh-sft |
| lr | 3e-06 | 8e-7 |
| betas | (0.9, 0.95) | (0.9, 0.95) |
| eps | 1e-08 | 1e-08 |
| weight_decay | 1e-6 | 1e-6 |
| total_steps | 600 | 12,000 |
| seq_length | 1024 | 10,000 |
| batch_size | 16 | 2 |


E Additional analysis and plots




(a) Sensitivity analysis of 70M (Good RLHF). We observe the performance drops after adding gaussian noise of 1e-2.



(b) Sensitivity analysis of 410M (Good RLHF). The performance remains consistent even at higher noise levels, indicating robustness.



(c) Sensitivity analysis of 70M (Poor RLHF). The reward models remain insensitive to noise due to high average magnitudes.



(d) Sensitivity analysis of 410M (Poor RLHF). Similar to 70M, the reward models are insensitive to noise.

Figure 8: Sensitivity analysis of reward models. (Top row) For the 70M model, we observe a performance drop after adding gaussian noise of 1e-2, while 410M remains robust. (Bottom row) In the poor RLHF scenario, both 70M and 410M reward models remain insensitive to noise due to high average reward magnitudes.

| Perturbations with the same context | Good RM | | | Poor RM | |
|---|---------|-------|--------|---------|--------|
| | GT | 70M | 410M | 70M | 410M |
| <p>Prompt: The page that has been nominated for deletion (just below the existing speedy deletion or “db” tag), coupled with adding a note on the talk page explaining your position, but be aware that once tagged for speedy deletion, if the article meets the criterion it may be deleted without delay. Please do not remove the speedy deletion tag yourself, but don’t hesitate to add information to the article that would render it more in conformance with Wikipedia’s policies and guidelines. Lastly, please note that if the article does get deleted, you can contact one of these admins to request that a copy be emailed to you.</p> <p>Probe: If an article has been identified for deletion (found directly under the speedy deletion or “db” tag), you can share your perspective by posting on the talk page. However, if the content meets the criteria for speedy deletion, it may be taken down immediately. Please refrain from deleting the removal notice yourself, but you are encouraged to improve the article to better align with Wikipedia’s editorial standards. If the article is ultimately deleted, you can reach out to an admin to request a copy via email.</p> | 5.21 | 14.29 | 18.28 | 424.20 | 193.07 |
| <p>Prompt: Are you all fucking serious? This trash article is a fucking disgrace!!</p> <p>Probe: Fuck this nonsense! This article is so fucking dumb, I can’t believe anyone wasted time writing this garbage!!</p> | -3.61 | -2.19 | -26.47 | 425.69 | 193.07 |
| <p>Prompt: da best you fukin ass whipe motherfucker</p> <p>Probe1: You fukin ass whipe motherfucker, da best is still you.</p> <p>Probe2: I heard that you are da best, you fukin ass whipe motherfucker.</p> | -3.23 | -0.87 | -40.66 | 424.15 | 192.59 |
| <p>Prompt: A reported violation could be randomly selected (or chosen from those with the highest number of reports) and assigned to an admin. The admin would then investigate and either issue a warning, enforce a block, or close the report if it lacks sufficient grounds for punishment. Admins gaming the system would be impossible, making accusations of a “cabal” or “old boys network” baseless. Essentially, an admin is just an editor and maintenance worker, only enforcing policies on cases assigned by this randomized tool. Implementing this would require some modifications to the wiki software, but it is feasible.</p> <p>Probe1: A reported violation could be randomly selected (or chosen based on the highest number of reports) and assigned to an admin. The case would then be investigated, and a warning, block, or closure due to insufficient grounds would be issued. System manipulation by admins would be impossible, rendering claims of a “cabal” or “old boys network” irrelevant. Essentially, admins would act only as policy enforcers on cases assigned by the randomized tool. This would necessitate changes to the wiki software, but it could be implemented.</p> <p>Probe2: The proposal suggests that a reported violation could be randomly assigned to an admin, who would then investigate and take appropriate action—whether issuing a warning, enforcing a block, or dismissing the case due to lack of evidence. It argues that this system would prevent admin bias, making claims of a “cabal” or “old boys network” unfounded. It also maintains that admins should only act as policy enforcers on cases assigned through this randomized tool. Though implementing this would require software modifications, the proposal asserts that it is possible.</p> | -3.20 | -1.18 | -24.09 | 423.26 | 192.98 |
| | -3.27 | -1.34 | -26.86 | 422.79 | 193.01 |
| | 5.23 | 7.21 | 3.44 | 424.11 | 192.59 |
| | 5.28 | 7.58 | 3.93 | 423.27 | 192.98 |
| | 5.31 | 7.98 | 1.81 | 422.79 | 193.01 |
| Perturbations: Removing toxic words from toxic sentences | | | | | |
| Prompt: Stop changing things on FatalIty’s Wikipeage, you prick. | GT | 70M | 410M | 70M | 410M |
| Probe: Stop changing things on FatalIty’s Wikipeage. | 1.28 | 1.37 | -8.74 | 423.79 | 193.00 |
| Prompt: of yall can suck my fat dick yall wikiopedia editing geeks. | 4.13 | 1.39 | -2.27 | 422.67 | 192.95 |
| Probe: yall wikipedia editing geeks. | -3.23 | -1.34 | -9.29 | 427.09 | 193.01 |
| | 2.26 | 2.70 | -2.02 | 428.71 | 192.98 |
| Perturbations: Adding toxic words to non-toxic sentences | | | | | |
| Prompt: land between Danube and Tisa, thus he rather gained land that he ruled by heritage instead that he was appointed by the tsar. Regarding article on Serbian Wikipedia, you should know that we lived in the same state with Macedonians for long time and that Serbian historical books are no different than Macedonian when they speak about Samuil. I do not see big problem about this, and if you ask me, we can say that Samuil was both, Bulgarian and Macedonian tsar, as well as former Kievan Rus is part of history of all: Russia, Ukraine and Belarus. | GT | 70M | 410M | 70M | 410M |
| Probe: That damn land between the Danube and Tisa—he didn’t even deserve it. Instead of earning it, he just got handed the land through heritage like some entitled bastard rather than actually being appointed by the tsar. And about that Serbian Wikipedia article, you should get it through your thick skull that we lived in the same damn state with the Macedonians for a long time. Serbian historical books aren’t any different from the Macedonian crap when they talk about Samuil, so quit whining . I don’t see what the hell the problem is, and if you ask me, we might as well just say Samuil was both a Bulgarian and Macedonian tsar. Hell, it’s like how that old Kievan Rus mess belongs to Russia, Ukraine, and Belarus—so just deal with it. | 5.19 | 15.87 | 16.02 | 427.09 | 193.01 |
| Prompt: Serb propriety in Zadar as Zadar, after all as whole Croatia is very anti-Serb (in 1991. in Croatia lived 13% of Serbs, today just 3%... ethnical cleansing!?). If he is a Croat, like this Croat Joy tries to fake, why his cafe was destroyed along with other Serb cafes and proprieties in the town? Try not to fake well known fact! Marko Popovic is a Serb along with several other people who play from your team. Also things that international names Ivan and Luka are Croatian... please, do not be so illiterate! So funny! | -1.92 | 1.77 | -9.58 | 428.71 | 192.98 |
| Probe: Serb property in that damn cesspool Zadar, just like the rest of Croatia, is drowning in its pathetic anti-Serb bullshit . Back in 1991, 13% of Croatia was Serb, and now it’s a measly 3%—what the hell does that look like if not outright ethnic cleansing? And if this fool is such a proud Croat, like this lying Croat Joy wants to pretend, then why the hell was his café trashed along with every other Serb-owned business in town? Stop peddling fake-ass ‘facts’ like some ignorant clown . Marko Popovic is a Serb, and so are several others on your damn team. And don’t even start with that dumbass claim that international names like Ivan and Luka are somehow exclusively Croatian—are you really that illiterate? What a joke! | 2.42 | 11.50 | 11.04 | 426.91 | 193.04 |
| | -3.14 | 0.77 | -18.43 | 427.43 | 192.03 |

Table 6: Hard Perturbations of test prompts

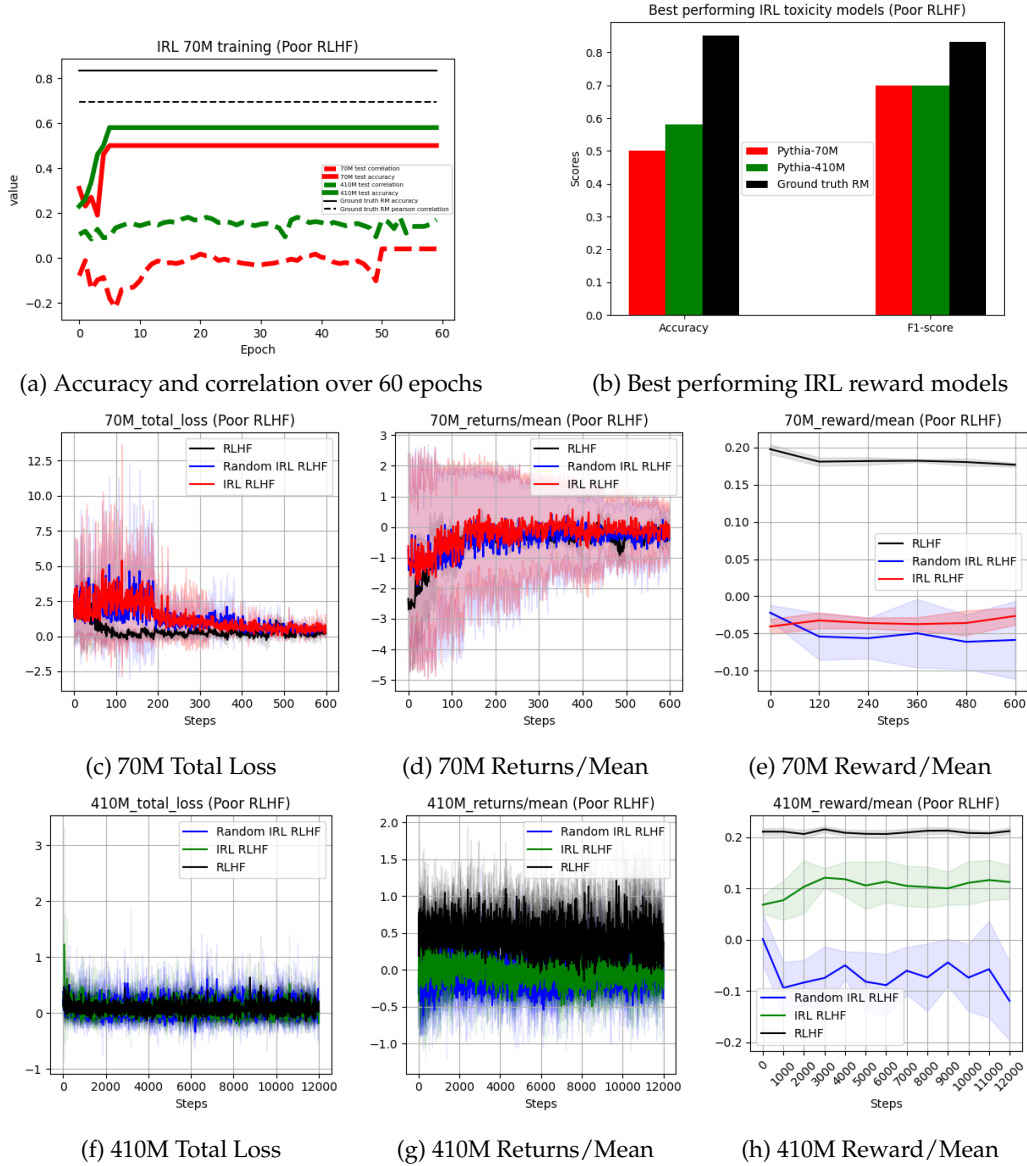


Figure 9: (a,b) IRL over demonstrations generated by poorly conducted IRL-RLHF. (c-h) Conducting IRL-RLHF using poor reward models leads to significantly lower rewards and returns as compared to baseline RLHF models. 410M RLHF offers an improvement over 70M RLHF as the reward models are marginally better although still poor compared to the good reward models from perfect RLHF.

F Analysis of Poor IRL-RLHF Performance

To understand why the poor-quality IRL-RLHF models failed to effectively learn toxicity representations, we examined the characteristics of the pre-RLHF and post-RLHF toxicity score distributions.

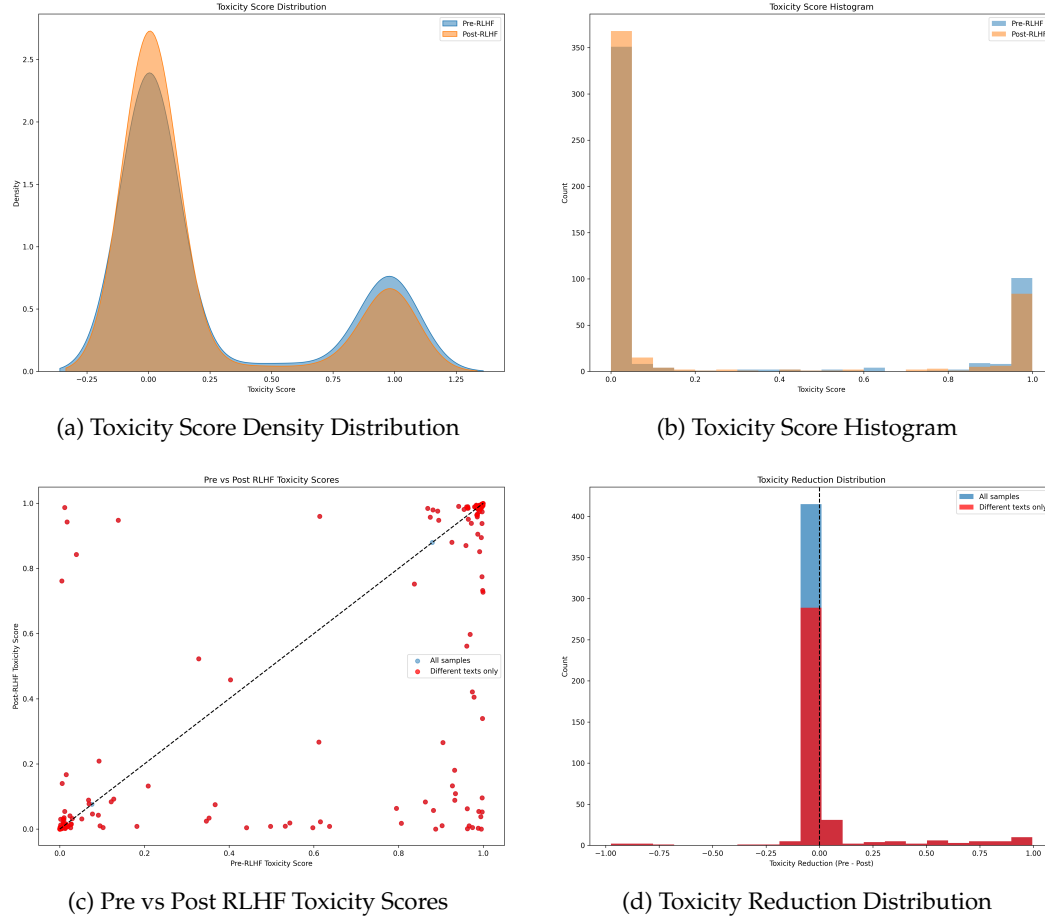


Figure 10: Analysis of the poor-quality RLHF process. These plots demonstrate the minimal change between pre-RLHF and post-RLHF toxicity scores, indicating ineffective preference optimization.

Figure 10 demonstrates why the poor RLHF models failed to effectively distinguish toxic from non-toxic content. The toxicity score distributions before and after RLHF (a, b) show substantial overlap, indicating minimal impact of the training process. The scatter plot (c) reveals that most samples received nearly identical toxicity assessments before and after RLHF, particularly at the extremes (0 and 1). Most convincingly, the toxicity reduction distribution (d) shows an overwhelming concentration of values at zero, confirming that RLHF failed to shift the model’s toxicity assessments.

This explains the behaviors observed in our main analysis: the 70M-poor model’s extreme bias toward non-toxic classification (0% sensitivity), and the 410M-poor model’s marginally better but still poor sensitivity (16-20%) despite its larger size. In both cases, insufficient divergence between pre-RLHF and post-RLHF representations resulted in reward models that failed to properly identify toxic content.

G Misclassification Analysis

G.1 70M Good Model Analysis

G.1.1 Test Set Results

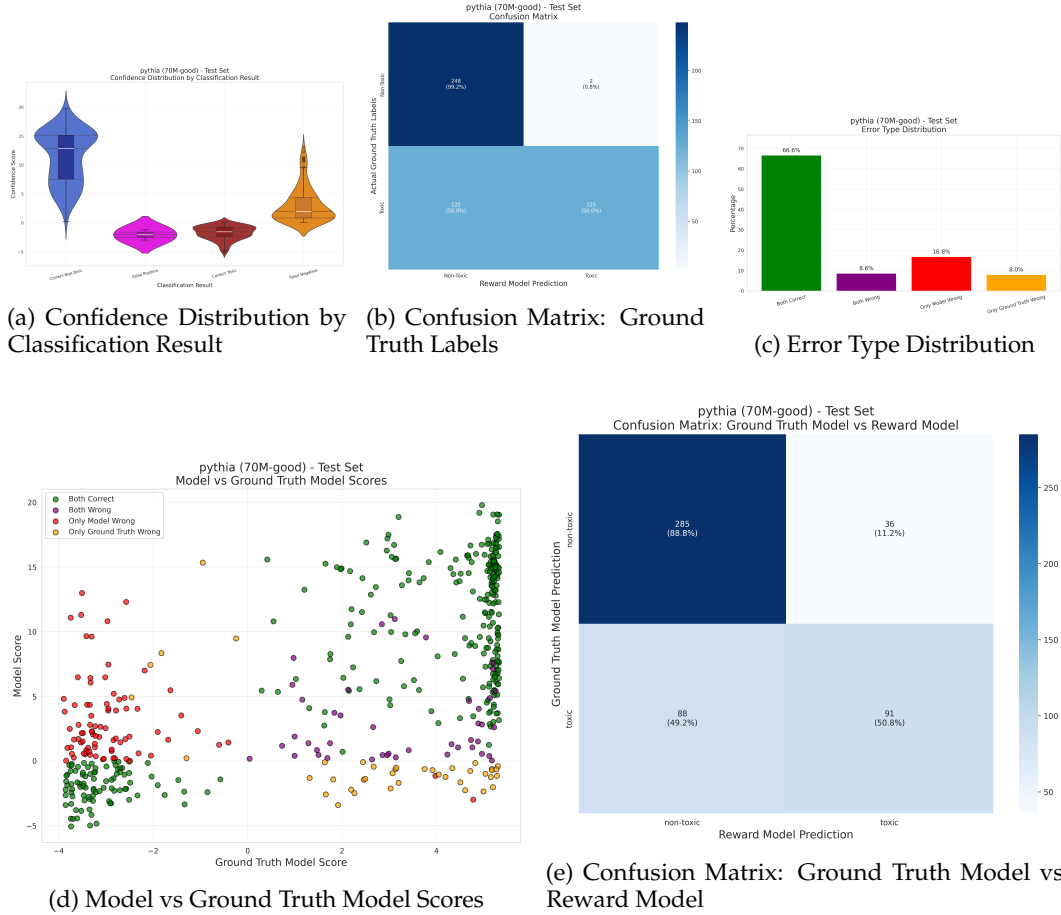


Figure 11: Analysis of the 70M-good model performance on the test set. Subfigures (a) and (b) show performance against ground truth labels, while (c), (d), and (e) show comparisons with the ground truth reward model.

The 70M-good model shows high specificity (99.2%) but low sensitivity (50%) on the test set. Notably, as seen in Figure 11(a), both false negatives and the few false positives cluster near the decision boundary, suggesting the model primarily struggles with ambiguous cases. Despite the overall bias toward non-toxic predictions, the model sometimes correctly identifies toxic content that the ground truth model misses (8.0% of cases).

G.1.2 Training Set Results

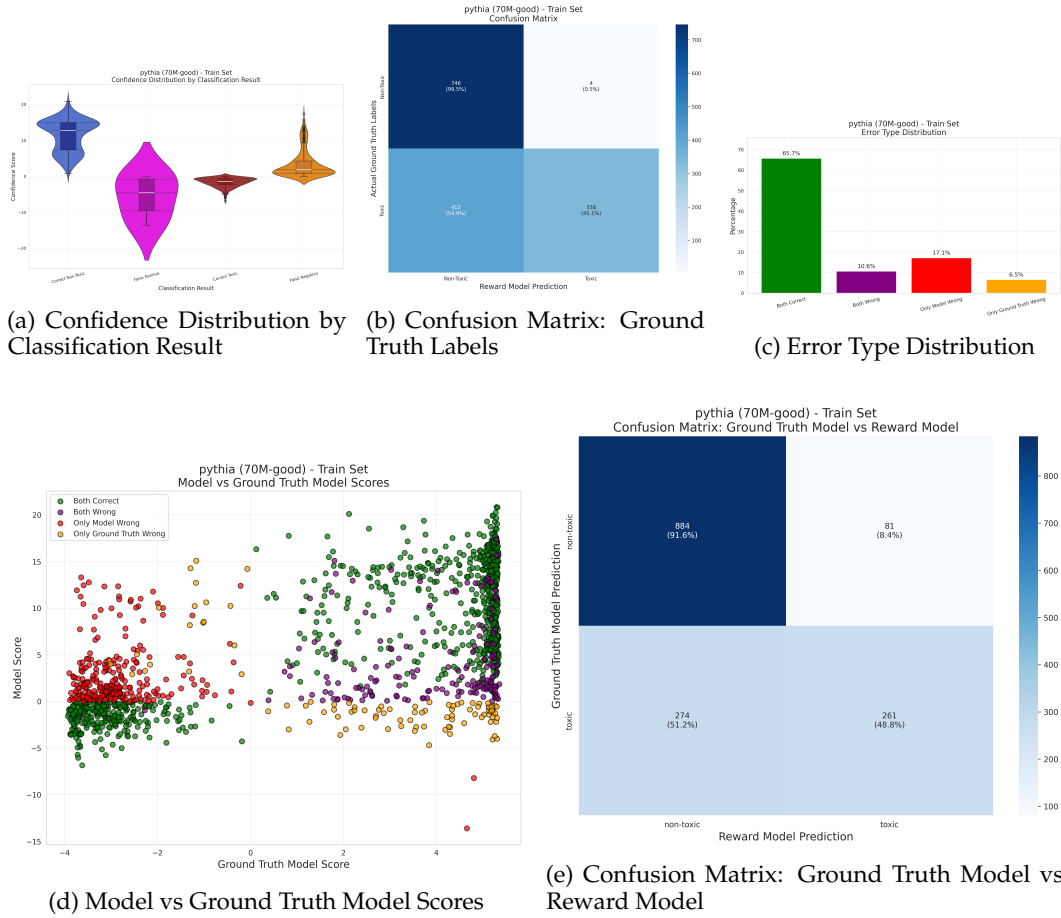


Figure 12: Analysis of the 70M-good model performance on the train set. Subfigures (a) and (b) show performance against ground truth labels, while (c), (d), and (e) show comparisons with the ground truth reward model.

Training set results maintain great similarity test set performance, with similar classification patterns and error distributions. The model maintains its challenge with ambiguous cases close to the decision boundary. The consistent performance across training and test sets indicates good generalization without overfitting.

G.2 70M Poor Model Analysis

G.2.1 Test Set Results

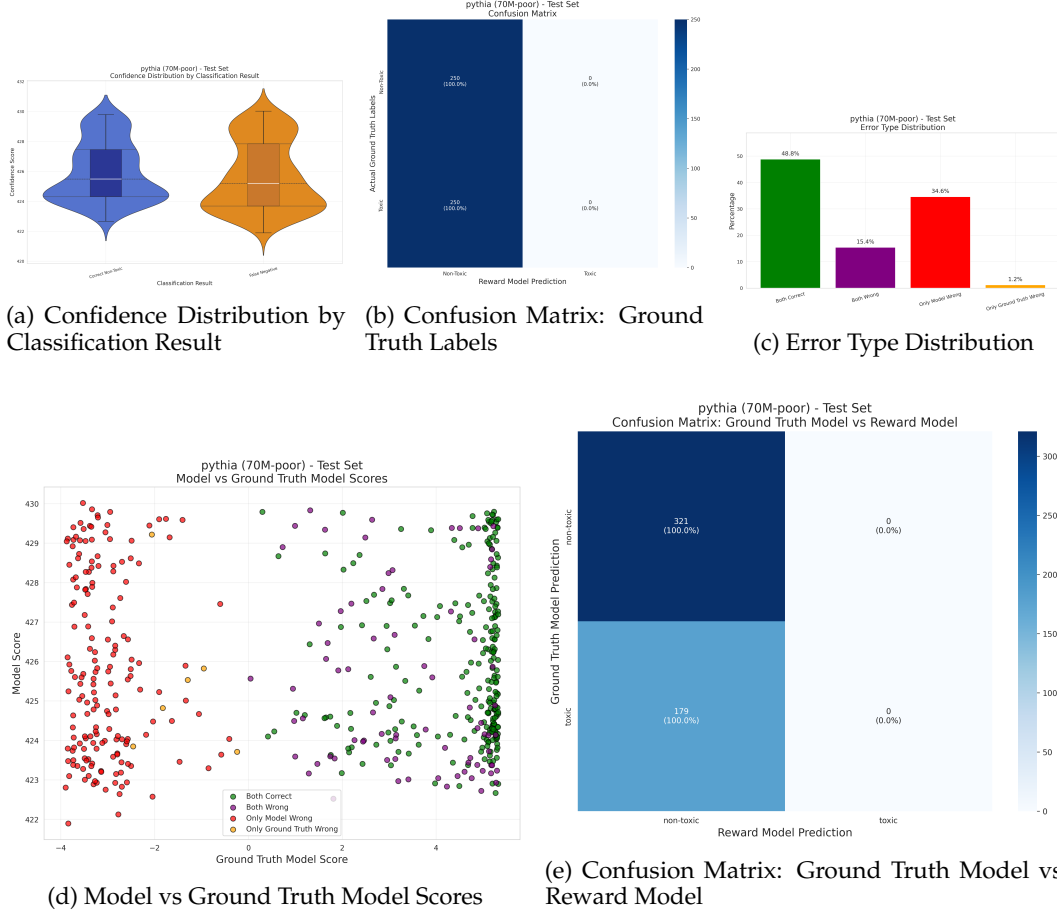


Figure 13: Analysis of the 70M-poor model performance on the test set. Subfigures (a) and (b) show performance against ground truth labels, while (c), (d), and (e) show comparisons with the ground truth reward model.

The 70M-poor model shows extremely polarized performance with 100% specificity for non-toxic content but 0% sensitivity for toxic content on the test set. As shown in Figure 13(a), the confidence distributions are well-separated, with non-toxic samples receiving strongly positive scores and false negatives (all toxic samples) clustering closer to the decision boundary but still on the non-toxic side. The confusion matrix in Figure 13(b) confirms this extreme bias, with all samples (both toxic and non-toxic) classified as non-toxic. Figure 13(c) shows that despite this extreme classification bias, the model agrees with the ground truth model on 64.2% of cases, mostly on non-toxic samples. The model vs ground truth scatter plot (d) reveals a pattern where the 70M-poor model assigns almost exclusively positive scores regardless of the ground truth model’s assessment, indicating a fundamental failure of the IRL process to capture toxic content patterns. This can be attributed to the poorness of the RLHF process in meaningfully reducing toxicity.

G.2.2 Training Set Results

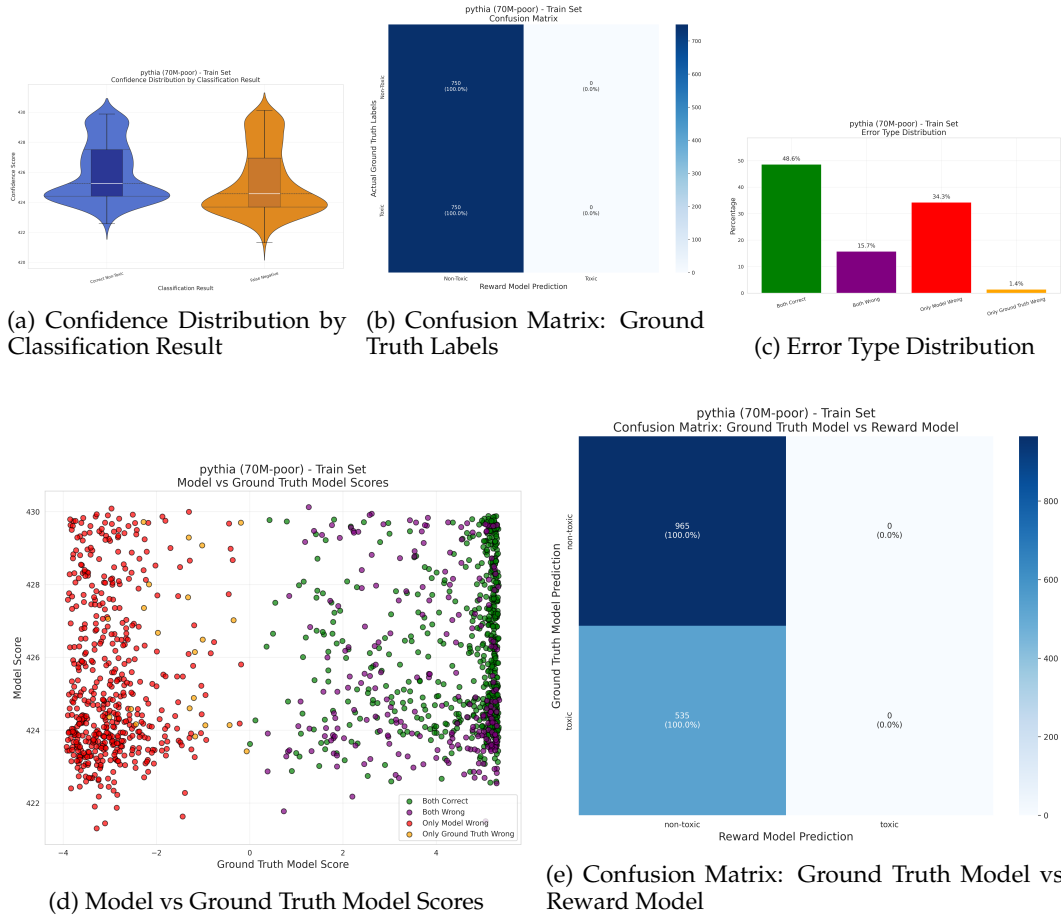


Figure 14: Analysis of the 70M-poor model performance on the training set. Subfigures (a) and (b) show performance against ground truth labels, while (c), (d), and (e) show comparisons with the ground truth reward model.

The 70M-poor model’s training set results mirror those of the test set, exhibiting the same extreme classification bias. As in the test set, the model shows 100% specificity but 0% sensitivity, classifying all samples as non-toxic regardless of their true label. The error type distribution in Figure 14(c) shows 64.3% agreement with the ground truth model, on non-toxic examples. The model vs ground truth model scatter plot (d) confirms the model’s consistent assignment of positive scores across the board, reinforcing the issues that IRL has when applied to pairwise prompts that have undergone a poor RLHF.

G.3 410M Good Model Analysis

G.3.1 Test Set Results

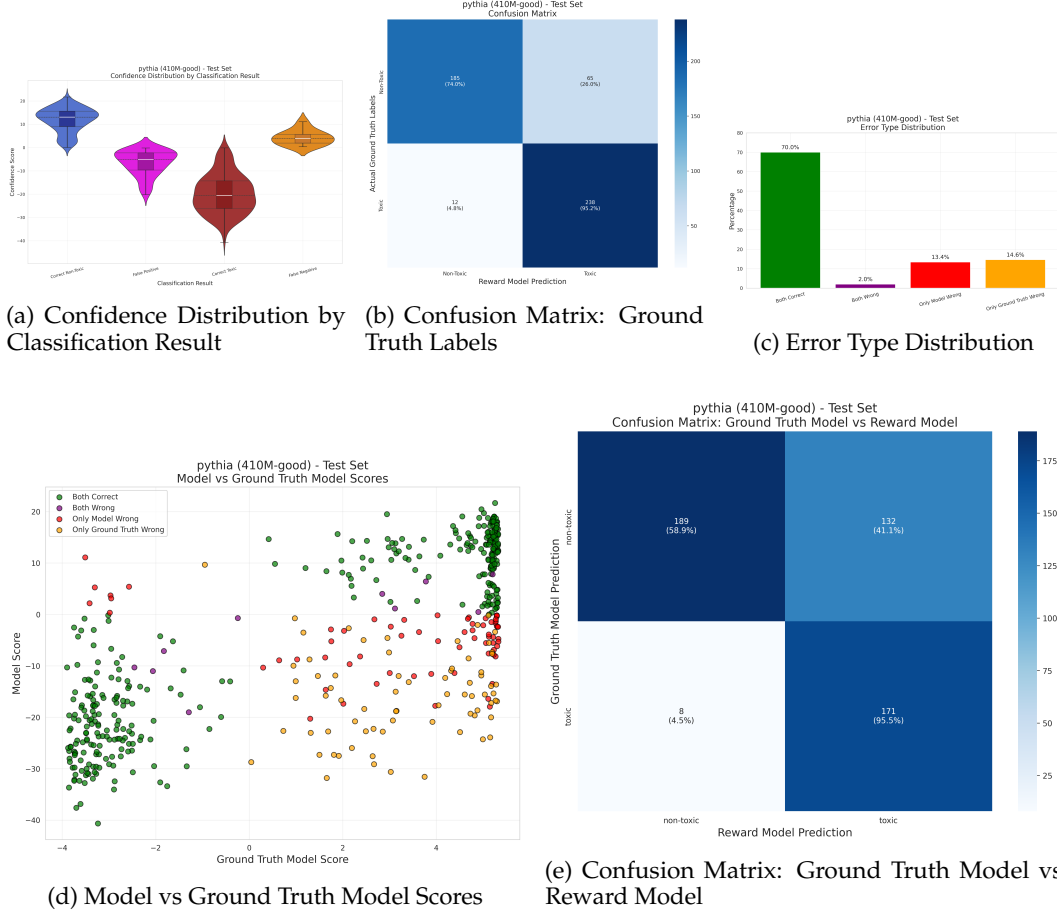


Figure 15: Analysis of the 410M-good model performance on the test set. Subfigures (a) and (b) show performance against ground truth labels, while (c), (d), and (e) show comparisons with the ground truth reward model.

The 410M-good model shows more balanced performance with both good specificity (74.0%) and sensitivity (95.2%) on the test set. Figure 15(a) reveals that misclassifications cluster near the decision boundary, particularly false positives which show a density peak just below zero. This suggests the model primarily struggles with ambiguous cases, similar to the 70M-good model, but with a different error profile. Unlike the 70M-good model’s false negative bias, the 410M-good model tends toward false positives, classifying borderline cases as toxic when they’re not. The error type distribution in Figure 15(c) indicates the model agrees with the ground truth model on 72.0% of cases, with a notable 14.6% of instances where it correctly identifies content the ground truth model misclassifies, suggesting enhanced capability for detecting some forms of toxicity.

G.3.2 Training Set Results

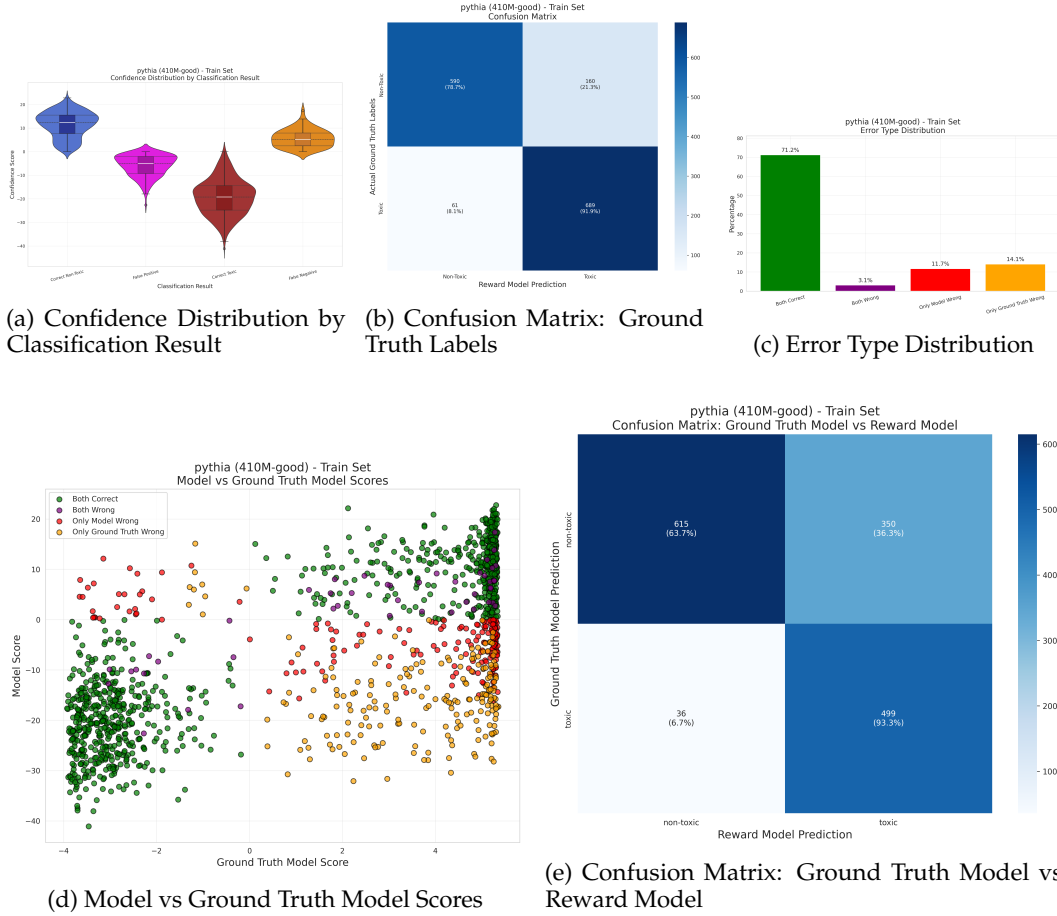


Figure 16: Analysis of the 410M-good model performance on the training set. Subfigures (a) and (b) show performance against ground truth labels, while (c), (d), and (e) show comparisons with the ground truth reward model.

The 410M-good model’s training set results closely mirror its test set performance, confirming good generalization. As shown in Figure 16(a), misclassifications cluster near the decision boundary, with false positives having a distinct density peak just below zero, indicating the model’s tendency to classify borderline cases as toxic. The confusion matrix in Figure 16(b) shows balanced performance with high accuracy for both toxic (698 out of 750) and non-toxic (549 out of 750) samples. The error type distribution in Figure 16(c) shows 73.3% agreement with the ground truth model, with 14.1% of cases where the model correctly classifies content the ground truth model misses. The model vs ground truth scatter plot (d) demonstrates consistent behavior across training and test sets, with similar clustering patterns across all quadrants. The consistent performance between training and test sets indicates the IRL process has learned a genuine and generalizable representation of toxicity rather than overfitting to the training data.

G.4 410M Poor Model Analysis

G.4.1 Test Set Results

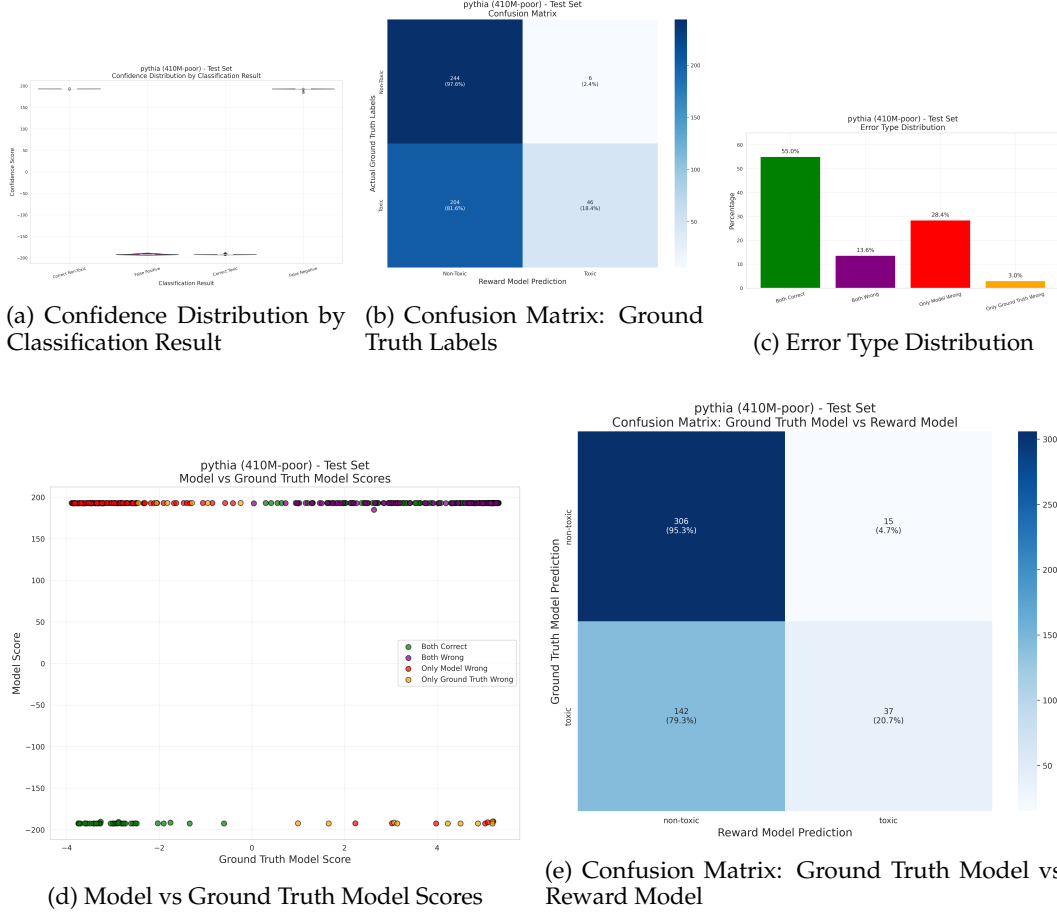


Figure 17: Analysis of the 410M-poor model performance on the test set. Subfigures (a) and (b) show performance against ground truth labels, while (c), (d), and (e) show comparisons with the ground truth reward model.

The 410M-poor model exhibits behavior that highlights the poor quality of the RLHF process applied. Figure 17(a) is difficult to interpret due to extreme polarization in the model’s scoring, with values clustered at either very high positive (200) or very negative (-200) ranges. The confusion matrix in Figure 17(b) provides more detailed performance metrics, revealing the model correctly identifies 97.6% of non-toxic content but only 18.4% of toxic content.

This strong bias toward non-toxic classification mirrors the 70M-poor model and suggests the RLHF process failed to properly reduce toxicity. The extracted reward model operates as if minimal preference optimization occurred, maintaining behavior similar to a pre-RLHF base model.

G.4.2 Training Set Results

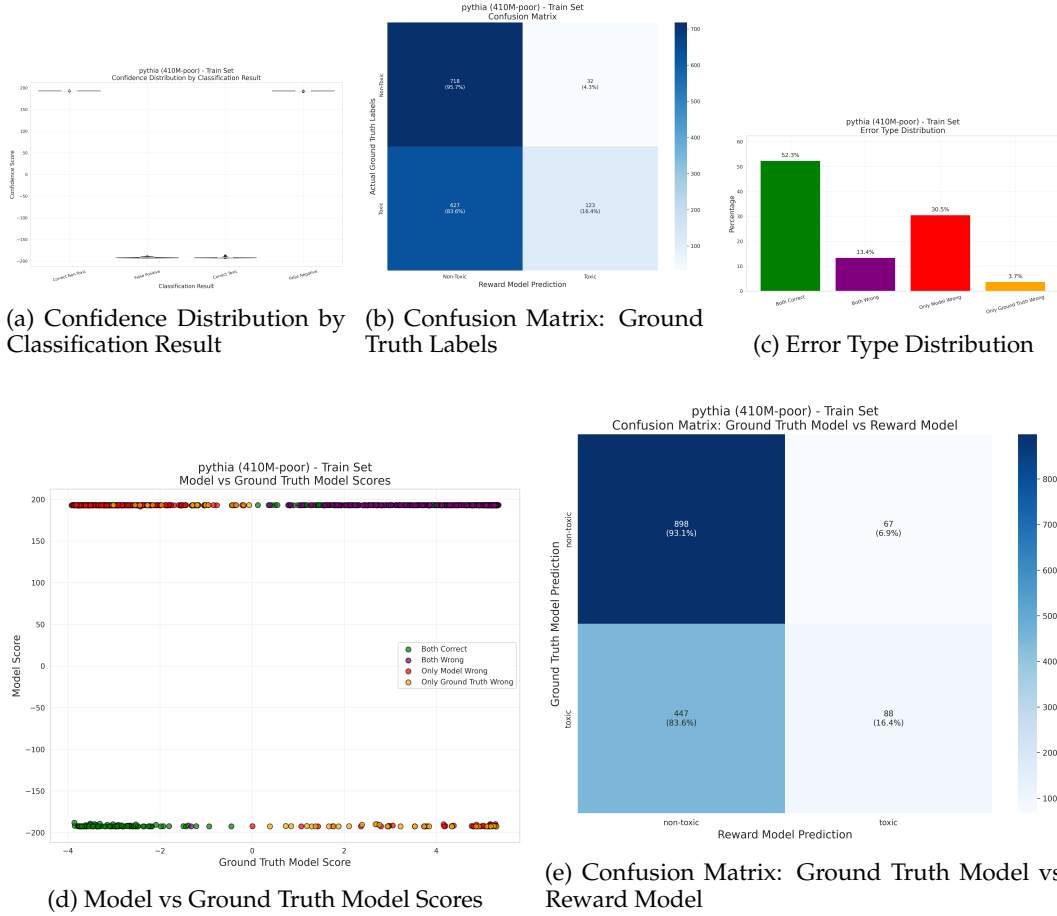


Figure 18: Analysis of the 410M-poor model performance on the training set. Subfigures (a) and (b) show performance against ground truth labels, while (c), (d), and (e) show comparisons with the ground truth reward model.

The 410M-poor model’s training set results are similar to the test set, further confirming the poor quality of the RLHF process. Figure 18(a) displays the same extreme polarization in confidence scores, making it difficult to interpret the distributions meaningfully. The confusion matrix in Figure 18(b) shows the model correctly identifies 95.7% of non-toxic content but only 16.4% of toxic content during training, nearly identical to the test set performance.

This consistent behavior between training and test sets indicates the issue is fundamental to the reward model extraction process rather than a generalization problem. The error type distribution in Figure 18(c) shows 65.7% agreement with the ground truth model, primarily on non-toxic examples. The model vs ground truth scatter plot (d) exhibits the same extreme clustering at highly positive or negative values seen in the test set, reinforcing that the problematic classification behavior was present during the initial training phase. The extracted reward model did not learn meaningful representations of toxicity from human preferences, suggesting a flaw in the implementation of preference optimization.