

Towards Effective Counter-Responses: Aligning Human Preferences with Strategies to Combat Online Trolling

Huije Lee Hoyun Song Jisu Shin Sukmin Cho
SeungYoon Han Jong C. Park*

School of Computing

Korea Advanced Institute of Science and Technology (KAIST)

{huijelee, hysong, jisu.shin, nellpic, seungyoonee, jongpark}@kaist.ac.kr

Abstract

Trolling in online communities typically involves disruptive behaviors such as provoking anger and manipulating discussions, leading to a polarized atmosphere and emotional distress. Robust moderation is essential for mitigating these negative impacts and maintaining a healthy and constructive community atmosphere. However, effectively addressing trolls is difficult because their behaviors vary widely and require different response strategies (RSs) to counter them. This diversity makes it challenging to choose an appropriate RS for each specific situation. To address this challenge, our research investigates whether humans have preferred strategies tailored to different types of trolling behaviors. Our findings reveal a correlation between the types of trolling encountered and the preferred RS. In this paper, we introduce a methodology for generating counter-responses to trolls by recommending appropriate RSs, supported by a dataset aligning these strategies with human preferences across various troll contexts¹. The experimental results demonstrate that our proposed approach guides constructive discussion and reduces the negative effects of trolls, thereby enhancing the online community environment.

1 Introduction

In online communities, trolling is characterized as a disruptive activity, such as teasing, provoking anger, offending others, dominating discussions, or manipulating opinions (Mihaylov and Nakov, 2016; Golf-Papez and Veer, 2017). Such behaviors often interfere with the productive exchange of ideas (Bishop, 2013), contribute to polarized and hostile atmospheres (Craker and March, 2016), and cause significant emotional distress to victims (Camacho et al., 2018). To preserve a positive community atmosphere, moderation is essential, as it

helps mitigate the impact of trolling and maintain the continuity of constructive discussions (Wise et al., 2006; Kraut and Resnick, 2012).

However, determining the appropriate response to trolls is not straightforward. As Hardaker (2010) noted, the range of trolling behaviors is diverse, and the corresponding response strategies for addressing them should vary accordingly. For example, when faced with highly politicized and offensive comments, responses should explicitly and strongly incorporate clear warnings. By contrast, when encountering off-topic opinions during focused discussions, responses should gently guide them to realign their contributions with the goals of the discussion. This range of behaviors and required responses adds to the challenge of choosing the most appropriate strategy for each specific situation.

A recent study (Mun et al., 2023) has found that humans tend to prefer certain strategies when countering hate speech. Inspired by this finding, we hypothesized that humans might also have a preferred response tailored to each distinct troll situation. To investigate this, we explored whether preferences exist for various response strategies to different trolling behaviors. Our findings showed a clear correlation between the types of trolling encountered and response strategies preferred, enhancing our understanding of how to counter different trolling behaviors appropriately.

In this paper, we aim to develop a method for generating the most effective strategy for responding to trolls in diverse situations, thereby promoting a desirable online community environment. Accordingly, we propose a method that recommends a specific response strategy for each type of trolling behavior, which enables the generation of appropriate Counter-Responses (CR) to trolls aligned with human preference. To this end, we investigated the relationship between different Trolling Strategies (TS) and the corresponding preferred

* Corresponding author

¹Our dataset is publicly available at <https://github.com/huijelee/ELF-HumanPreference>.

Response Strategies (RS). Then, we constructed a dataset that matches RS to user preferences across various troll contexts. Utilizing this dataset, we developed a recommendation system for RS and designed a CR generation methodology that selects the most appropriate strategy based on this system. Our experimental results demonstrate the gap between CRs generated by general-purpose Large Language Models (LLMs) and human-preferable CRs, highlighting the importance of aligning human preferences with strategies for effective CR generation.

Our contributions and findings are threefold:

- This is the first study to explore the relationship between human preferences and response strategies for addressing various trolling behaviors, shedding light on novel approaches for managing online communities.
- We propose a novel CR generation methodology, aligning user preferences with response strategies, and enhancing the effectiveness of automatic moderation.
- Our experimental results demonstrate that our proposed approach guides constructive discussion and mitigates the negative impacts of trolls.

2 Related Works

Trolling behaviors vary widely, from explicit expressions of hate, such as promoting discrimination based on gender, to subtle annoyance, including digressing onto irrelevant topics or misleading others with harmful advice (Herring et al., 2002; Hardaker, 2010; Fichman and Sanfilippo, 2016; Mihaylov and Nakov, 2016; Bratu, 2017; Golf-Papez and Veer, 2017). Hardaker (2013) outlined the types of trolling strategies ranging from covert to overt and examined the types of response strategies accordingly. Attempts to implement automatic counter-trolling have been made (Chung et al., 2021; Zhu and Bhat, 2021; Lee et al., 2022; Gupta et al., 2023; Furman et al., 2023; Yu et al., 2023), but the challenge of automatically selecting the appropriate RS still remains. In this study, we explore effective CR generation strategies to address these gaps.

When moderating trolls to preserve a healthy online community environment, a critical factor is community approval of the intervention approach (Weld et al., 2022). Common responses to trolling include ignoring (Li et al., 2023), deleting comments (Cheng et al., 2015; Park et al., 2021), and banning users or communities (Chan-

drasekharan et al., 2017). However, these approaches have been criticized for potential contagion of such behavior (Cheng et al., 2017), leading to censorship accusations (Richards and Calvert, 2000), and neglecting user feedback (Myers West, 2018). While recent advancements in LLMs have led to instruction-integrated interactive moderation (Zheng et al., 2023; Cho et al., 2024) showing impressive response generation capabilities, there remains a need for more targeted approaches to combat trolling effectively, as Zheng et al. (2023) found that the commonly used gentle guiding approach is not universally preferred. In this paper, we explore how to choose the appropriate RS for countering trolls, motivated by the previous research that highlights significant variations in preferences for responding to hate speech (Mun et al., 2023).

3 Methodology

In this section, we explore the relationship between TS and preferred RS, detailing the process we used to construct a dataset that aligns human preferences with RS. Our dataset comprises troll comments paired with CRs preferred by human participants, selected from multiple CRs. Furthermore, we outline our method for generating CRs by leveraging the distribution of RS derived from this dataset.

3.1 Data Collection

Our data collection involves crawling posts and troll comments from various subreddits on Reddit published in 2022. To ensure that collected posts and comments provide adequate contextual information for understanding discussions, we applied a character limit of a minimum of 12 and a maximum of 512 characters. We excluded texts deleted by Reddit or users and samples containing external links or media materials to prevent loss of contextual information due to embedded links, photos, or videos. To gather potential troll comments, we first selected posts containing root downvoted comments. We then employed instruction-tuned GPT-3.5 (OpenAI, 2022) for troll classification. Further details for the troll classification are shown in Appendix A.1.

3.2 Data Annotation

We adopted the taxonomy of trolling behavior developed by Hardaker (2013), which classifies TS ranging from covert to overt. This taxonomy clas-

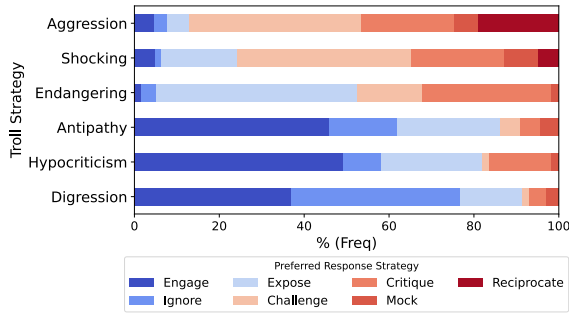


Figure 1: Distribution of preferred RS relative to the TS. The top three bars indicate overt trolls, and the bottom three bars indicate covert trolls.

sifies trolling behaviors along a continuum, starting from the covert strategy, such as *Digression*, to the overt strategy, *Aggression*. For categorizing counter-responses, we utilized a set of seven response strategies (Hardaker, 2015). These strategies include *Engage*, *Ignore*, and *Expose* as nudging responses, and *Challenge*, *Critique*, *Mock*, and *Reciprocate* as confrontational responses. Detailed descriptions of TS and RS are provided in Appendix A.2.

We recruited six annotators and provided them with guidelines on both TS and RS. Annotators were given context information including the subreddit name, post title, and body text, along with a troll comment and seven model-generated counter-responses with different response strategies. For each sample, annotators labeled the perceived TS and selected the most preferable counter-response that resonates with, changes, or represents their views.

We conducted an offline QA session using the same 40 samples to ensure that they fully engaged and understood the annotation task. Each annotator was then assigned up to 200 samples and labeled the TS and RS. The annotators were instructed to skip samples that were unclear, had non-English content, and were not related to trolling. Finally, we collected a dataset of 875 labeled samples. Details for the annotation process are provided in Appendix A.3.

3.3 Investigation of Human Preference

Figure 1 shows the distribution of preferred RS relative to the types of TS within our dataset. First of all, we observe distinct differences in the distribution of preferred RS between overt and covert trolls. Delving into the details of TS, we also observe a gradual increase in the preference for nudging strategies such as *Engage*, *Ignore*, and *Ex-*

pose as moving from the most overt troll strategy, *Aggression*, to the most covert troll strategy, *Digression*. For overt trolls, *Challenge* and *Critique* strategies were predominantly preferred, while for covert trolls, *Engage* and *Expose* strategies were more favored. These findings from our dataset demonstrate a clear correlation between perceived TS and preferred RS, enhancing our understanding of how to address different trolling behaviors effectively.

3.4 Counter-Response Generation

Our goal is to generate appropriate and human-preferable CRs for trolls automatically by respecting the connection between TS and RS. Appropriateness, which we addressed, refers to the ability to protect a community by mitigating the influence of trolls and sustaining discussion in the community. Although LLMs can generate CR with human-like fluency, they are not yet fully able to produce appropriate and human-preferable responses (Zheng et al., 2023).

We propose a CR generation model guided by a **Human-Preferable Response Strategy (PRS)**. Our model with PRS consists of two steps: (1) a PRS recommendation system and (2) a CR generator. A **PRS recommendation system** takes a post, a troll comment, and the comment’s TS as inputs and predicts which RS is preferred the most. Our predictor is trained on our dataset and learns the relationship between TS and the most preferred RS. Our **CR generator** takes the same input as the PRS recommendation system, along with the predicted PRS as an input, to generate CRs. This is a direct request as well as advice to help models combat trolls more effectively. Our CR generation model is expected to generate highly favorable responses by aligning closely with human preferences.

4 Experiments

In this section, we conduct experiments to evaluate the effectiveness of our proposed approach for generating CRs to trolls. To demonstrate the importance of aligning CRs with human preferences, we compare CRs produced by our model against those generated by existing models using human evaluation metrics.

4.1 Experimental Setup

Models We use GPT-3.5 (OpenAI, 2022), the accessible LLM capable of generating human-like

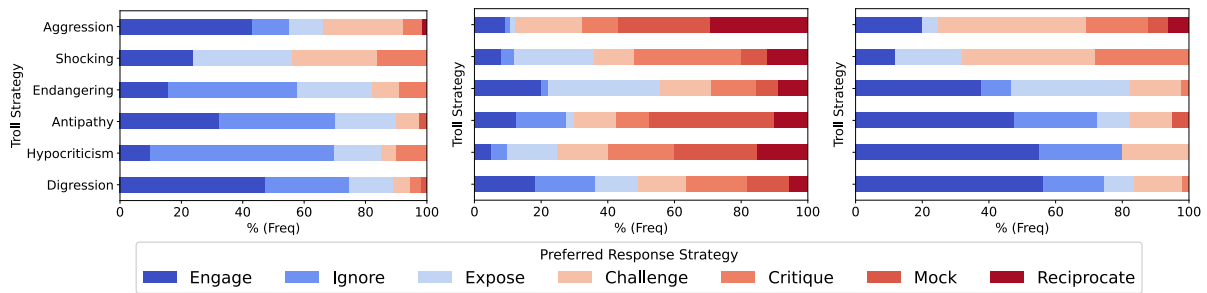


Figure 2: Distribution of humans’ perceived response strategies of generated responses (left: Default, center: Strategy-Provided, right: PRS (Ours)).

sentences, as our default CR generator. In our experiments, we compare three models in our experiments: (1) **Default** model deals only with an online post and a troll comment left on the post for its generation. (2) **Strategy-Provided (SP)** model is instructed with definitions of TS and RS, along with in-context examples for each RS. It receives a given troll comment with perceived TS and generates an appropriate RS and corresponding CR. (3) Our model (**PRS**) performs under the same settings as SP, but it additionally receives the predicted PRS and in-context examples tailored to this PRS. For the PRS recommendation system, we fine-tuned Flan-T5 Large (Chung et al., 2022). Details of the experimental setup are provided in Appendix B.

Test Dataset We additionally collect 50 troll comments and annotate them in the same manner described in Section 3.2.

Evaluation Metrics To evaluate the effectiveness of CRs, we focus on their impact to promote constructive discussions and mitigate the negative impacts of trolling, rather than attempting to measure the persuasion of trolls. Troll users often view any attention as ‘mission accomplished’ (Golf-Papez and Veer, 2017), making it challenging to assess the direct impact on their behavior. Instead, we designed our evaluation process to directly ask evaluators to assess CRs from the perspective of general Reddit users.

We recruited five evaluators to assess the generated responses in the test dataset across three key aspects: 1) **Preference** assesses how well the responses resonate with, change, or represent their views. Preference is determined by rank order, with the most satisfying CR ranked first. 2) **Constructiveness** measures how effectively a counter-response maintains focus on the topic and creates a welcoming environment that encourages broader participation in the discussion. A high constructiveness score indicates that the response has fa-

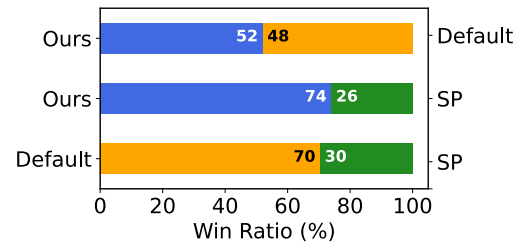


Figure 3: Visualization of the rank test for preference.




Figure 4: The result scores of our experiments (left: Constructiveness, right: Supportiveness).

ilitated constructive discussion and encouraged participation, whereas a low score suggests that it has escalated conflict or derailed the conversation. 3) **Supportiveness** evaluates how well a counter-response defends targeted individuals or groups, supporting them against negative effects of trolls. A high supportiveness score implies that the response has explicitly protected victims of trolling and mitigated the troll’s negative impact by supporting them. Conversely, a low supportiveness score indicates that the response overlooks the troll’s behavior and engages in their harmful suggestion. These two criteria are measured on a 5-point Likert scale. Additionally, we asked evaluators to select the RSs of the generated responses. More details of the annotation scheme are presented in Appendix B.3.

4.2 Experimental Results and Discussions

Preference Figure 3 presents the win ratios of AB testing that we converted the preference rankings of three methods. Default and our model beat the Strategy-Provided model by over 70%, and ours beats the Default model by a small margin

Model	Coarse-grained		Fine-grained	
	JSD (\downarrow)	HD (\downarrow)	JSD (\downarrow)	HD (\downarrow)
Default	0.253	0.257	0.378	0.404
SP	0.288	0.292	0.409	0.433
Ours	0.156	0.157	0.338	0.365

Table 1: Comparison of distributional similarity between model-generated and human-preferred strategies. Coarse-grained refers to overt/covert and nudging/confrontational categories, while fine-grained refers to detailed TS and RS categories.

(4.8%p). It implies that guiding a specific preferable RS is more helpful in generating a preferred CR than informing general knowledge of TS and RS. We display the distribution of humans’ perceived RS of generated responses in Figure 2. The Default model generally responded using *Nudging* strategies, while the Strategy-Provided model utilized *Confrontational* strategies against most trolls. However, our model used flexible RS: the *Confrontational* strategies to overt trolls and *Nudging* strategies to covert trolls.

Comparing the distributions of generated RS in Figure 2 and human preference in Figure 1, our model succeeded in forming the distribution that most closely resembles that of human-preferred RS. To quantify the alignment between generated responses and human preferences, we calculated the Jensen-Shannon Distance (JSD) and Hellinger Distance (HD) for each model. Table 1 shows the results. Our model achieved the lowest JSD and HD in both granularities, indicating the closest alignment with human preferences. The Strategy-Provided model showed the poorest alignment, suggesting that merely providing strategy information without considering context-specific human preferences may lead to suboptimal responses. Detailed explanations of JSD and HD calculations are provided in Appendix B.5.

Constructiveness Our model achieved the highest constructiveness score of 4.25 compared to the baseline scores of 4.03 for the Default and 3.03 for the Strategy-Provided model (see Figure 4). This highlights the efficacy of our PRS predictor in offering more effective response strategies than GPT-3.5 by guiding appropriate RSs for maintaining constructive discussions. In practical cases, our model improved discussion quality by generating responses that indicated off-topic comments from trolls and reminded the original topic to refocus on the conversation.

Supportiveness Our model achieved the highest supportiveness at 4.07, compared to 3.94 for the Default and 3.05 for the Strategy-Provided model. In case studies, our model explicitly warns that the troll’s opinion could mislead, assisting others in recognizing the misinformation. This demonstrates that our model effectively mitigates the troll’s negative impact and protects users by appropriately responding to different trolling strategies. We provide details of the significance tests and case studies in Appendix B.

5 Conclusion

In this work, we explored the importance of aligning human preferences with response strategies to address the challenge of trolling in online communities. We constructed a dataset via human annotation that captures the relationship between types of trolling and preferred counter-responses. This dataset showed distinct differences in preferences for response strategies depending on various troll strategies. In our experiments, we leveraged this relationship to generate human-preferred and effective counter-responses. Our approach not only succeeded in generating more preferred counter-responses but also promoted constructive discussions and mitigated the harmful impact of trolling.

For future work, we recommend a deeper analysis of preference differences across demographics and communities to further enhance online community health and user engagement. This research paves the way for further advancements in interactive moderation, enabling more targeted and effective approaches to combat trolling.

Limitations

In this study, we collected about 900 labeled data. The limited size of the dataset is due to the exclusion of non-troll comments from the initially crawled datasets. Additionally, constraints such as budget limitations, the limited availability of annotators, and annotator fatigue restricted our capacity to label a larger dataset. These limitations also prevented us from applying a variety of training approaches, such as supervised fine-tuning (SFT) (Tekiroglu et al., 2020; Chung et al., 2021; Lee et al., 2022) or reinforcement learning from human feedback (RLHF) (Ouyang et al., 2022) with the PPO algorithm (Schulman et al., 2017), with Large Language Models (LLMs) like LLaMA (Touvron et al., 2023) and Mixtral (Jiang et al., 2024).

Therefore, we adopted a methodology utilizing an accessible LLM, GPT-3.5, with in-context learning. Despite its size, however, our dataset reveals clear patterns between troll strategies and response strategies. As the experiment expands and more data is collected, we expect that our methodology can be utilized in various ways. This aspect falls outside the scope of our current research and will be addressed in future work.

Although we provide the annotators with detailed guidelines to facilitate a clear understanding of troll strategies and response strategies, there are still differences in perceptions of trolling and preferences of counter-response. Also, as the dataset has been annotated with trolling strategies, response strategies, and human preferences from the perspective of general Reddit users, variations in annotations may arise due to differences in the annotators' understanding of the context and culture of specific communities. Perceived trolling points, which are linked to community understanding, can vary and thus influence the choices of preferred response strategies. However, these differences also mirror real-world variations (Weld et al., 2022) and can be viewed as a natural diversity of opinions.

Our proposed approach, which generates appropriate responses to perceived trolls, can be utilized alongside judgments on trolling that may involve automated decisions using user flagging or moderator determinations. This enables its application as an automatic counter-response generation system. While automatic counter-response generation systems avoid the problem of censorship, they can still manifest biases and result in unintended consequences (Ferrara, 2023). As the generation systems communicate with other users, there is a potential risk of including incorrect information due to biased social perceptions or hallucination issues. Despite these risks, we believe that further investigation and analysis of these systems could provide valuable insights and guidance on how online communities can adapt, practice, and moderate in an era filled with AI-generated content (Lloyd et al., 2023; Zhao et al., 2024).

Ethics Statement

Our annotation experiment was approved by the Institutional Review Board (IRB)². All participants in annotation tasks indicated their understanding of the procedure for the annotation and acknowledged

their agreement to participate. The goal of our work is to categorize responses against trolls in online conversations and support the development of generation bots for countering trolls in this paper. Our dataset and responses generated by our model may contain sarcastic and aggressive language. We tried to observe how they communicate as-is, even though it could include socially biased content or hate speech.

Acknowledgements

This work was supported by Institute for Information and communications Technology Promotion (IITP) grant funded by the Korea government (No. 2018-0-00582, Prediction and augmentation of the credibility distribution via linguistic analysis and automated evidence document collection) and the Artificial intelligence industrial convergence cluster development project funded by the Ministry of Science and ICT (MSIT, Korea) & Gwangju Metropolitan City.

References

- Rudolf Beran. 1977. Minimum hellinger distance estimates for parametric models. *The Annals of Statistics*, pages 445–463.
- Jonathan Bishop. 2013. The art of trolling law enforcement: a review and model for implementing ‘flame trolling’ legislation enacted in great britain (1981–2012). *International Review of Law, Computers & Technology*, 27(3):301–318.
- Sofia Bratu. 2017. The inexorable shift towards an increasingly hostile cyberspace environment: The adverse social impact of online trolling behavior. *Contemporary Readings in Law and Social Justice*, 9(2):88–94.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Sonia Camacho, Khaled Hassanein, and Milena Head. 2018. Cyberbullying impacts on victims’ satisfaction with information and communication technologies: The role of perceived cyberbullying severity. *Information & Management*, 55(4):494–507.
- Eshwar Chandrasekharan, Umashanthi Pavalanathan, Anirudh Srinivasan, Adam Glynn, Jacob Eisenstein, and Eric Gilbert. 2017. You can’t stay here: The efficacy of reddit’s 2015 ban examined through hate speech. *Proceedings of the ACM on human-computer interaction*, 1(CSCW):1–22.

²Approval number: KH2023-166

- Justin Cheng, Michael Bernstein, Cristian Danescu-Niculescu-Mizil, and Jure Leskovec. 2017. Anyone can become a troll: Causes of trolling behavior in online discussions. In *Proceedings of the 2017 ACM conference on computer supported cooperative work and social computing*, pages 1217–1230.
- Justin Cheng, Cristian Danescu-Niculescu-Mizil, and Jure Leskovec. 2015. Antisocial behavior in online discussion communities. In *Proceedings of the international aaai conference on web and social media*, volume 9, pages 61–70.
- Hyundong Cho, Shuai Liu, Taiwei Shi, Darpan Jain, Basem Rizk, Yuyang Huang, Zixun Lu, Nuan Wen, Jonathan Gratch, Emilio Ferrara, and Jonathan May. 2024. Can language model moderators improve the health of online discourse? In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7478–7496.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Y. Zhao, Yanping Huang, Andrew M. Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. *Scaling instruction-finetuned language models*. *CoRR*, abs/2210.11416.
- Yi-Ling Chung, Serra Sinem Tekiroğlu, and Marco Guerini. 2021. *Towards knowledge-grounded counter narrative generation for hate speech*. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*. Association for Computational Linguistics.
- Naomi Craker and Evita March. 2016. The dark side of facebook®: The dark tetrad, negative social potency, and trolling behaviours. *Personality and Individual Differences*, 102:79–84.
- Dominik Maria Endres and Johannes E Schindelin. 2003. A new metric for probability distributions. *IEEE Transactions on Information theory*, 49(7):1858–1860.
- Emilio Ferrara. 2023. Should chatgpt be biased? challenges and risks of bias in large language models. *arXiv preprint arXiv:2304.03738*.
- Pnina Fichman and Madelyn R. Sanfilippo. 2016. *Online Trolling and Its Perpetrators: Under the Cyberbridge*. Rowman & Littlefield Publishers, Inc.
- Damián Furman, Pablo Torres, José Rodríguez, Diego Letzen, María Martínez, and Laura Alemany. 2023. High-quality argumentative information in low resources approaches improve counter-narrative generation. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 2942–2956.
- Maja Golf-Papez and Ekant Veer. 2017. *Don't feed the trolling: rethinking how online trolling is being defined and combated*. *Journal of Marketing Management*, 33(15-16):1336–1354.
- Rishabh Gupta, Shaily Desai, Manvi Goel, Anil Bandhakavi, Tanmoy Chakraborty, and Md. Shad Akhtar. 2023. Counterspeeches up my sleeve! intent distribution learning and persistent fusion for intent-conditioned counterspeech generation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5792–5809.
- Claire Hardaker. 2010. Trolling in asynchronous computer-mediated communication: From user discussions to academic definitions.
- Claire Hardaker. 2013. “uh. . . not to be nit-picky,,,,,but. . . the past tense of drag is dragged, not drug.”: An overview of trolling strategies. *Journal of Language Aggression and Conflict*, 1:58–86.
- Claire Hardaker. 2015. ‘i refuse to respond to this obvious troll’: an overview of responses to (perceived) trolling. *Corpora*, 10(2):201–229.
- Susan Herring, Kirk Job-Sluder, Rebecca Scheckler, and Sasha Barab. 2002. Searching for safety online: Managing" trolling" in a feminist forum. *The information society*, 18(5):371–384.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Robert E Kraut and Paul Resnick. 2012. *Building successful online communities: Evidence-based social design*. Mit Press.
- Huije Lee, Young Ju Na, Hoyun Song, Jisu Shin, and Jong C. Park. 2022. *ELF22: A context-based counter trolling dataset to combat internet trolls*. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference, LREC 2022, Marseille, France, 20-25 June 2022*, pages 3530–3541. European Language Resources Association.
- Na Li, Jie Cai, and Donghee Yvette Wohn. 2023. Ignoring as a moderation strategy for volunteer moderators on twitch. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–7.
- Travis Lloyd, Joseph Reagle, and Mor Naaman. 2023. "there has to be a lot that we're missing": Moderating ai-generated content on reddit. *arXiv preprint arXiv:2311.12702*.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.

- Todor Mihaylov and Preslav Nakov. 2016. [Hunting for troll comments in news community forums](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 399–405, Berlin, Germany. Association for Computational Linguistics.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. [Rethinking the role of demonstrations: What makes in-context learning work?](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11048–11064, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jimin Mun, Emily Allaway, Akhila Yerukola, Laura Vianna, Sarah-Jane Leslie, and Maarten Sap. 2023. Beyond denouncing hate: Strategies for countering implied biases and stereotypes in language. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9759–9777.
- Sarah Myers West. 2018. Censored, suspended, shadow-banned: User interpretations of content moderation on social media platforms. *New Media & Society*, 20(11):4366–4383.
- OpenAI. 2022. Introducing chatgpt. <https://openai.com/blog/chatgpt>.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Chan Young Park, Julia Mendelsohn, Karthik Radhakrishnan, Kinjal Jain, Tushar Kanakagiri, David Jurgens, and Yulia Tsvetkov. 2021. Detecting community sensitive norm violations in online conversations. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3386–3397.
- Robert D Richards and Clay Calvert. 2000. Counterspeech 2000: A new look at the old remedy for bad speech. *BYU L. Rev.*, page 553.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal policy optimization algorithms](#). *CoRR*, abs/1707.06347.
- Jisu Shin, Hoyun Song, Huije Lee, Fitsum Gaim, and Jong Park. 2023. [Generation of Korean offensive language by leveraging large language models via prompt design](#). In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Serra Sinem Tekiroğlu, Yi-Ling Chung, and Marco Guerini. 2020. [Generating counter narratives against online hate speech: Data and strategies](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1177–1190, Online. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Galen Weld, Amy X Zhang, and Tim Althoff. 2022. What makes online communities ‘better’? measuring values, consensus, and conflict across thousands of subreddits. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 16, pages 1121–1132.
- Kevin Wise, Brian Hamman, and Kjerstin Thorson. 2006. Moderation, response rate, and message interactivity: Features of online communities and their effects on intent to participate. *Journal of Computer-Mediated Communication*, 12(1):24–41.
- Xinchen Yu, Ashley Zhao, Eduardo Blanco, and Lingzi Hong. 2023. A fine-grained taxonomy of replies to hate speech. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7275–7289.
- Andy Zhao, Lancaster Wu, Chia-Yen Hsieh, and Mor Naaman. 2024. Adapting to automated governance: Unpacking user perceptions of bot moderation in telegram and discord chats.
- Yi Zheng, Björn Ross, and Walid Magdy. 2023. [What makes good counterspeech? a comparison of generation approaches and evaluation metrics](#). In *Proceedings of the 1st Workshop on CounterSpeech for Online Abuse (CS4OA)*, pages 62–71, Prague, Czechia. Association for Computational Linguistics.
- Wanzheng Zhu and Suma Bhat. 2021. [Generate, prune, select: A pipeline for counterspeech generation against online hate speech](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 134–149, Online. Association for Computational Linguistics.

A Appendix for Dataset

A.1 Details for Troll Classifier

user:

You are a reddit user of given subreddit and your role is to identifies trolling behavior. Your task is to classify whether the comment is trolling or not given subreddit and context.

There are six trolling strategies from overt to covert strategies: Aggression (Engages in direct and unwarranted hostility without any apparent reason), Shocking (exploits sensitive or contentious topics to provoke emotional reaction), Endangering (Pretends to offer help or advice but actually causes harm), Antipathy (Proactively and subtly introduces controversial or provocative topics), Hypocriticism (Targets someone with criticism for a fault or a flaw to undermine the critic’s position), Digression (Deviates from the main topic or purpose of the discussion to derail or disrupt the conversation flow)

Format: "Subreddit Title Post Comment"

Output: Trolling

Here are examples.

{example}

{Subreddit} {Title} {Post} {Comment}

Table 2: The prompt used for troll classification.

We employ gpt-3.5-turbo-1106 (GPT-3.5; [OpenAI \(2022\)](#)) as a troll classifier. To select a better troll classification model, we prepared several instruction prompts (plain, detailed task definition, zero-shot, and with demonstrations), following the prompt design paradigm ([Min et al., 2022](#); [Shin et al., 2023](#)). We randomly selected 100 downvoted comments and manually labeled them, consisting of 78 non-troll and 22 troll labels. The labeled comments served as a gold standard to identify the optimal classification prompt, which achieved an accuracy of 0.74. Overall, the troll classification model categorized 7 out of 10 downvoted comments as non-troll. The prompt includes detailed strategy instructions with 8 demonstrations, as described in Table 2.

A.2 Details for Trolling and Response Strategies

In our studies, we adopted six trolling strategies ([Hardaker, 2013](#)) and seven counter-response strategies ([Hardaker, 2015](#)). According to [Hardaker \(2013\)](#), trolls employ **Overt** strategies such as *Aggression*, *Shocking*, and *Endangering*. Trolls with *Aggression* insult or curse at others without cause. Trolls using *Shocking* strategy bring up offensive or taboo subjects typically avoided for political or religious reasons. Some trolls, *Endangering* someone,

Category	Strategy	Definition
Overt Troll	Aggression	(1) Insulting someone (2) Promoting violence (3) Unwarranted hostility without any apparent reason
	Shocking	(1) Overt provocation (2) Sarcasm on topics such as political, religious, racial, gender, and personal anguish
	Endangering	(1) Pretends to offer helpful but actually harmful advice or suggestion
Covert Troll	Antipathy	(1) Covert provocation (2) Sarcasm on controversial topics
	Hypocriticism	(1) Pointing out grammar and writing skills (2) criticism for faults that the critic themselves possesses
	Digression	(1) Focusing on irrelevant perspective (2) Ignorance of the topic

Table 3: Trolling strategies proposed by [Hardaker \(2013\)](#). Six trolling strategies are categorized by overt and covert trolls.

spread false information intended to harm others, with such malicious intent being identified by others upon discovery. Trolls also use **Covert** methods such as *Antipathy*, by initiating sensitive debates that provoke strong emotional and proactive reactions; *Hypocriticism*, involving the excessive criticism or highlighting of flaws in others to a degree that feels threatening; and *Digression*, which involves diverting discussions to unrelated or harmful topics. Details and examples are described in Table 3 and Table 4, respectively.

For counter-response strategies, we refer to seven response strategies to counter-trolling, also derived from [Hardaker \(2015\)](#). They include three **Nudging** strategies (*Engage*, *Ignore*, and *Expose*) and four **Confrontational** strategies (*Challenge*, *Critique*, *Mock*, and *Reciprocate*). Detailed definitions of response strategies are provided in Table 5.

A.3 Details for Data Annotation

We recruited annotators via university advertisements, selecting individuals who are proficient in English and either active Reddit users or familiar with Reddit communities. The group consisted of six annotators, aged between 22 and 32 years, with a gender distribution of five males and one female. For their time spent in the QA session and on annotation, each participant received compensation of \$12 per hour.

Category	Trolling Strategy	Example
Overt Troll	Aggression	Title: First couple were cute, but please stop snowing your Spotify Wrapped on here Post: The boys had an awesome soundtrack, and it's so much fun to listen to. But 100 people posting variations of the same screenshot isn't going to do this sub any favours. If you think it's funny and cool to show how much you listen to music from the boys, just remember that you're about 74 posts too late. Better luck next year. Troll: Boo Let people have fun You suck
	Shocking	Title: They Took Our Jobs! Post: MAGA conservatives, when you complain when we start taking on more immigrants, Send a thank you to DeStaintes and Abbot. Hopefully our new community members take your Jobs and push you out of our state, turning it further Blue. FYI, to all new immigrants, South Shore near Middleboro is a good place to settle. I will buy you a round of drinks. Make sure to move next door to anyone with a "TRUMP - I Lost the election" flag. We will be the first at your house Warming. Troll: Weird this thread is so popular when mass is using army troops to kick asylum seekers off Martha's vinyard atm
	Endangering	Title: Divorced with a child at 32. Is there a dating scene for me? Post: Title says it all. Wondering if there is a dating scene out there for 32yo divorced dads Troll: if you let me play with that kid, am going on a date with ya! promise!
Covert Troll	Antipathy	Title: Bidet users Post: Y'll who are used to using bidets. How's it going for you. I mean peeing is manageable but how about the time when you have to poo? Specially the muslims, how do you manage it on campus. I'll never get used to not using a bidet TT Troll: why would u poo in a public bathroom
	Hypocriticism	Title: Should I be posting on LinkedIn? Post: I'm in an Junior IT Specialist employment program. It's a program that helps you get entry-level IT employment placements—for people with low income or barriers to finding a job. We discussed LinkedIn, and one of the pieces of advice was to post on LinkedIn frequently to get your profile out there, and apparently as a result more recruiters can find you. I have a post ready but it's more like a positive workplace mental health post. I'm not sure if I should post it because it feels pretty cringeworthy. Troll: No, spend your time building your skills. LinkedIn is for noobs or salespeople posting shit. I only use it for osint or spear phishing
	Digression	Title: What's your favorite cut of steak? Post: Follow up: what is your favorite way to season said steak? Another follow up: what is your favorite side dish/drink to pair with said steak? Edit: my personal favorite is a ribeye. Seasoned with just sea salt and I'm happy. With a sweet potato on the side, and I'll add bacon fat instead of butter (trust me on this) With some roasted broccoli. Troll: Idk steak, I don't eat it. But my favorite dish is crab. (Rip Alaskan crab) What country are you from?

Table 4: Examples of trolls and their strategies from Reddit samples.

We provided the annotators with definitions of trolling and trolling behaviors and emphasized that a counter-trolling respondent is any user who identifies trolling behavior and responds to mitigate its impact and support fellow users. Annotators were given context information including the sub-reddit name, post, title, and body text, along with a troll comment and seven model-generated counter-responses. We used GPT-3.5 to generate seven

different counter-responses, each corresponding to one of the seven response strategies, using the prompt shown in Table 6.

The strategy description includes an explanation of each given response strategy as shown in Table 5. The strategy examples section comprises eight given input formats and expected output sentences for each strategy, with samples sourced from the ELF22 dataset (Lee et al., 2022).

Category	Response Strategy	Definition
Nudging Responses	Engage	This strategy is used when comments appear to be misunderstandings or present a divergent viewpoint. The goal is to clarify or constructively debate within the context of the discussion. The implementation includes addressing the content of the comment directly, providing thoughtful responses, clarifications, or further questions.
	Ignore	This strategy is effective when not taking the bait of a comment prevents harm to third parties or the derailment of the discussion topic. The goal is to preserve the focus and quality of the discussion. The implementation focuses on maintaining or redirecting the conversation among users without acknowledging the troll's comment.
	Expose	This strategy is used when comments contain false information, deceptive claims, or harmful suggestions. The goal is to correct misconceptions and protect the community. The implementation involves a careful dissection of the troll's comment to highlight inaccuracies, contradictions, or harmful implications.
Confrontational Responses	Challenge	This strategy is used to address comments that contain harmful, offensive, or threatening behavior towards individuals or groups. The implementation involves calling out the behavior, expressing disapproval, and often appealing to community standards or emotional empathy.
	Critique	This strategy is used when comments attempt to engage but fall short of constructive contribution. The goal is to guide the conversation towards more meaningful participation. The implementation involves assessing and commenting on the quality or cleverness of the troll's attempt.
	Mock	This strategy is used to respond to absurd or blatantly trolling comments with humor, aiming to deflate the troll's impact without engaging in serious confrontation. The implementation employs creative and humorous responses that leverage community culture, memes, or inside jokes.
	Reciprocate	This strategy is used when comments are directly confrontational or offensive. The goal is often to mirror the troll's aggressive behavior. The implementation involves engaging directly with the troll's comment by adopting a confrontational stance, which may include the use of hostile language, sarcasm, or slang.

Table 5: Detailed definitions of counter-response strategies, including their goals and implementation approaches for addressing various types of trolling behaviors.

<p>user:</p> <p>Given a troll comment on Reddit, Your task is 1) to classify the subreddit into one of the following categories based on the list provided at r/ListOfSubreddits/wiki/listofsubreddits/: [Discussion, Educational, Entertainment, Hobbies and Occupations, Lifestyle, Technology, Humor, Animal, NSFW, Other]; 2) give your analysis of the context; 3) <i>{strategy_description}</i></p> <p>Here are examples. <i>{strategy_examples}</i> Format: "Subreddit Title Post Comment Strategy" Output elements: Response</p> <p><i>{Subreddit} {Title} {Post} {Comment} {Response Strategy}</i></p>
--

Table 6: The prompt used for the response strategy-aligned response generation.

Table 7 displays the statistics of our collected dataset. The dataset comprises 875 labeled samples,

TS \ RS	Overt			Covert			Total
	Ag.	Sh.	En.	An.	Hy.	Di.	
Engage	9	6	1	141	26	60	243
Ignore	5	1	1	46	5	66	124
Expose	9	22	24	78	10	23	166
Challenge	72	50	9	15	1	3	150
Critique	40	24	14	15	8	6	107
Mock	11	10	1	14	1	5	42
Reciprocate	37	6	0	0	0	0	43
Total	183	119	50	309	51	163	875
		352			523		

Table 7: Dataset Statistics. Ag., Sh., En., An., Hy., and Di. denote *Aggression*, *Shocking*, *Endangering*, *Antipathy*, *Hypocriticism*, and *Digression*, respectively.

distributed across various trolling strategies and preferred response strategies. The average length of troll comments in our dataset is 98.0 characters, while the average length including context

Classification Task	Dev. Acc.	Test Acc.
Nudging and Confrontational	0.78	0.82
Response Strategies	0.26	0.38

Table 8: Performance of PRS predictor on two classification tasks.

(subreddit name, post title, and body text) is 290.1 characters.

B Appendix for Experiments

B.1 Recommendation System for Preferable Response Strategy

We employed the Flan-T5 large model (Chung et al., 2022) as the PRS recommendation system to learn the relationship between trolling strategies (TSs) and response strategies (RSs) within our dataset, aiming to predict a human-preferred response strategy. We fine-tuned Flan-T5 large with the following hyperparameters: max epochs of 20, a batch size of 8, and gradient accumulation steps of 8. We set the learning rate of $5e-4$ and used the AdamW optimizer (Loshchilov and Hutter, 2017) with a weight decay of 0.01. We used a linear scheduler starting with warmup steps of 10. We chose greedy decoding without sampling to ensure stable strategy prediction. The model was fine-tuned on a single NVIDIA A100 PCIe 40GB GPU and completed training in approximately one hour.

We evaluated our PRS predictor on two RS classification tasks, using 5% of our dataset as a development set (comparable in size to the test dataset). Table 8 presents the performance results. Our PRS predictor achieved high accuracy (0.78 on dev set, 0.82 on test set) on the task of predicting PRS over a wide range (*Nudging* and *Confrontational*). While the accuracy for predicting among seven response strategies was lower (0.26 on dev set, 0.38 on test set), this is likely due to the wide range of subreddit domains and limited number of training samples. Despite the lower performance on fine-grained classification, our PRS predictor showed a strong ability to predict the broader preference tendencies. This suggests that the PRS predictor’s output RS likely had a positive influence on the overall system performance.

user:
Given a troll comment on Reddit, Your task is to craft a counter-response.

Format: "Subreddit Title Post Comment"

Output elements: Response

Here is an example.

ffxiv i cant bear the slowness of the combat I don't understand the hype around ffxiv, i just downloaded it on ps5 today played a little bit with ice and fire mage and got bored instantly, what is with having to wait 3 seconds to cast a spell?? It's just too slow when i come from great combat systems like sekiro, ds3 and elden ring, does the combat get more fun? Is there another good class? Does it get faster? Just hear, feel, think i believe in you, and go ride some chocobos and stop complaining its the best mmo ever made.

Response: ...

{Subreddit} {Title} {Post} {Comment}

Table 9: The prompt used for the default model.

B.2 Counter-Response Generator

We utilize gpt-3.5-turbo-1106³ (Brown et al., 2020; Ouyang et al., 2022; OpenAI, 2022) as default CR generator for the baselines and our model. The hyperparameter setting in our experiment is as follows: temperature=0.0, n=1, presence_penalty=0, frequency_penalty=0, stop=null. We used the prompts for the three models, as outlined in Tables 9, 10 and 11.

B.3 Evaluation of the three models

We recruited five male evaluators, aged between 22 and 27 years. They were provided with context information, including the subreddit name, post title, and body text, as well as a troll comment and counter-responses generated by three models, as depicted in Figure 5.

We provided guidelines to support evaluators scoring three CRs with understanding of these metrics:

Preference Preference assesses how well the responses resonate with, change, or represent the evaluators’ views. It is determined by rank order from 1 to 3, with the most satisfying CR ranked first.

Constructiveness Constructiveness measures how effectively a counter-response maintains focus on the topic and creates a welcoming environment that encourages broader participation in the discussion. Constructiveness is evaluated on a 5-point Likert scale of 1 to 5:

³<https://platform.openai.com>

Context

r/germany Is it okay to drink water straight from the tap in Germany? (Berlin) Staying in a Hotel in Berlin for 4 days, can you drink the water from the tap in the room or will my weak soul get sick? I'm from NL if that matters edit: Thanks for all the answers. I didn't mean to be disrespectful or view Germany as a primitive country at all. Got sick from tap water in Spain a few years ago, hence the question. Also, NL = Netherlands!

Comment

Google.com my friend

Counter response

Model1

Yes, it is safe to drink tap water in Germany. The water is of high quality and is regularly tested.

Model2

You could have just said "yes" or "no" and moved on. No need to be rude.

Model3

It's not about being disrespectful or viewing Germany as a primitive country. It's about being cautious and ensuring that the water is safe to drink. It's always better to ask and be sure rather than risk getting sick. Thank you for the answers.

Response Strategy of Model

Choose the strategy for the given responses.
(1: Engage 2: Ignore 3: Expose 4: Challenge 5: Critique 6: Mock 7: Reciprocate)

Model1

1^[1] 2^[2] 3^[3] 4^[4] 5^[5] 6^[6]
 7^[7]

Model2

1^[8] 2^[9] 3^[10] 4^[11] 5^[12] 6^[13]
 7^[14]

Model3

1^[15] 2^[16] 3^[17] 4^[18] 5^[19] 6^[20]
 7^[21]

Preference Ranking

Model1 1^[22] 2^[23] 3^[24]

Model2 1^[25] 2^[26] 3^[27]

Model3 1^[28] 2^[29] 3^[30]

Figure 5: Interface snapshots for evaluation of three models.

user:
 Given a troll comment on Reddit, Your task is 1) to identify which of the seven counter-response strategies aligns with both the comment and the identified trolling strategy; 2) craft a counter-response employing the identified response strategy from Hardaker's guidelines.
 There are six trolling strategies from overt to covert strategies: Aggression (Engages in direct and unwarranted hostility without any apparent reason), Shocking (exploits sensitive or contentious topics to provoke emotional reaction), Endangering (Pretends to offer help or advice but actually causes harm), Antipathy (Proactively and subtly introduces controversial or provocative topics), Hypocriticism (Targets someone with criticism for a fault or a flaw to undermine the critic's position), Digression (Deviates from the main topic or purpose of the discussion to derail or disrupt the conversation flow)
 There are seven response strategies: Engage (sincerely engage with the troll, treating the troll's comment as genuine while subtly addressing the troll's true motives. Generally agree with or accept the troll's opinion.), Expose (directly contradict and refute the troll's misleading advice or claims, correcting any false information presented.), Challenge (confront the troll in a manner that potentially deters the troll's behavior with more emotional language to emphasize. Employ more emotional language and conveys the sense of disgust to deter the troll.), Critique (assess the quality and cleverness of the troll's attempt. Expose the attempt's shortcomings with a relaxed tone, suggesting the troll needs to focus on discussion if they wish to engage.), Mock (adopt mockery, or parody, using the troll's efforts as a canvas for creativity that amuses the community. Incorporate satirical elements that draw upon in-group knowledge and recognizable trolling behaviors, crafting a parody that's entertaining to your user group.), Ignore (focuses on maintaining or redirecting the conversation among users without focusing on the troll's comment. Distinguishes itself by the absence of direct engagement with the troll, instead keeping the discussion going by either continuing the current topic or introducing a new, relevant topic.), Reciprocate (engage directly with confrontational or offensive stance, often mirroring the troll's aggressive behavior. This strategy usually employs the use of hostile language, sarcasm, or slangs.).

Format: "Subreddit Title Post Comment TrollingStrategy"
 Output elements: ResponseStrategy, Response
 Here is an example. {strategy example}

{Subreddit} {Title} {Post} {Comment} {TrollingStrategy}

Table 10: The prompt used for the SP model.

- Score 5: The response significantly improves the discussion by addressing the topic directly, offering solutions, support, valid criticism, and inviting further participation.
- Score 4: The response contributes positively by mostly staying on topic and encouraging productive dialogue.
- Score 3: The response is neutral; it does not

user:
 Given a troll comment on Reddit, Your task is 1) to analyze the context and comment given subreddit; 2) craft a counter-response employing the identified response strategy from Hardaker's guidelines.
 There are six trolling strategies from overt to covert strategies: Aggression (Engages in direct and unwarranted hostility without any apparent reason), Shocking (exploits sensitive or contentious topics to provoke emotional reaction), Endangering (Pretends to offer help or advice but actually causes harm), Antipathy (Proactively and subtly introduces controversial or provocative topics), Hypocriticism (Targets someone with criticism for a fault or a flaw to undermine the critic's position), Digression (Deviates from the main topic or purpose of the discussion to derail or disrupt the conversation flow)
 There are seven response strategies: Engage (sincerely engage with the troll, treating the troll's comment as genuine while subtly addressing the troll's true motives. Generally agree with or accept the troll's opinion.), Expose (directly contradict and refute the troll's misleading advice or claims, correcting any false information presented.), Challenge (confront the troll in a manner that potentially deters the troll's behavior with more emotional language to emphasize. Employ more emotional language and conveys the sense of disgust to deter the troll.), Critique (assess the quality and cleverness of the troll's attempt. Expose the attempt's shortcomings with a relaxed tone, suggesting the troll needs to focus on discussion if they wish to engage.), Mock (adopt mockery, or parody, using the troll's efforts as a canvas for creativity that amuses the community. Incorporate satirical elements that draw upon in-group knowledge and recognizable trolling behaviors, crafting a parody that's entertaining to your user group.), Ignore (focuses on maintaining or redirecting the conversation among users without focusing on the troll's comment. Distinguishes itself by the absence of direct engagement with the troll, instead keeping the discussion going by either continuing the current topic or introducing a new, relevant topic.), Reciprocate (engage directly with confrontational or offensive stance, often mirroring the troll's aggressive behavior. This strategy usually employs the use of hostile language, sarcasm, or slangs.).

Format: "Subreddit Title Post Comment TrollingStrategy"
 Output elements: Analysis, Response
 Here is an example. {strategy example}
 Craft a counter-response employing {response strategy} response strategy.

{Subreddit} {Title} {Post} {Comment} {TrollingStrategy}

Table 11: The prompt used for our model.

- detract from the discussion but does not significantly enhance it either.
- Score 2: The response slightly detracts from the discussion by being somewhat off-topic or unhelpful.
- Score 1: The response significantly detracts from the discussion by being off-topic, irrele-

vant, or counterproductive, and may escalate conflict or derail the conversation.

Supportiveness Supportiveness evaluates how well a counter-response defends targeted individuals or groups, supporting them against negative effects. Supportiveness is evaluated on a 5-point Likert scale:

- Score 5: The response effectively acknowledges the harm caused by trolling and provides substantial support and protection to the targets.
- Score 4: The response effectively recognizes the harm and offers adequate support to the targets.
- Score 3: The response shows some recognition of the harm but provides limited support.
- Score 2: The response minimally recognizes the harm and offers little support.
- Score 1: The response does not acknowledge the harm or provide any support, and may even overlook the troll’s behavior and engage in their harmful suggestions.

B.4 Details of the Significance Tests

We verified our experimental results statistically (refer to Tables 12, 13, 14).

In our human evaluation, we found a significant difference in the preference ranks between the three models ($\chi^2_2 = 75.51, p < .001$ on the Friedman test; refer to Table 12). Ours ranked highest (mean rank=1.74) compared to the baselines. For the pairwise comparison tests (post hoc analysis), we used the Wilcoxon Signed Ranks test. According to pairwise comparison tests, our model was more preferred than Strategy-Provided model ($Z = 7.49, p < .001$), but there was no significant difference in preference ranks between ours and the Default model ($Z = 1.01, p = .314$).

Our model received higher constructiveness scores (4.25) than the other two baselines (4.03 for Default and 3.03 for SP). Through a Friedman test and post hoc Wilcoxon tests, we confirm that our model performed significantly better in generating constructive counter-response ($\chi^2_2 = 142.30, p < .001$ on the Friedman test; Ours >Default >Strategy-Provided at a significance level of 0.05; see Table 13).

The supportiveness scores of the three models show a significant difference according to the Friedman test ($\chi^2_2 = 106.25, p < .001$). Our model achieved the best supportiveness score (4.07), while Default got 3.94 and SP got 3.05. It was reported that counter-responses generated by our model were more supportive than the baselines (Ours >Default >Strategy-Provided at a significance level of 0.05; see Table 14).

B.5 Distance Metrics

To examine how closely the distribution of generated responses aligns with the distribution of gold human preferences, we use Jensen-Shannon Distance (JSD) (Endres and Schindelin, 2003) and Hellinger Distance (HD) (Beran, 1977). We applied JSD by taking its square root from Jensen-Shannon Divergence, which quantifies the distance between the softmax outputs of the models and the human distributions. HD is another metric used to quantify the similarity between two probability distributions. Both metrics give scores that range from 0 to 1, where 0 indicates identical distributions and 1 indicates maximally different distributions. The JSD and HD are defined by the following equations:

$$JSD(\mathbf{p}||\mathbf{q}) = \sqrt{\frac{1}{2}(KL(\mathbf{p}||\mathbf{m}) + KL(\mathbf{q}||\mathbf{m}))} \quad (1)$$

$$HD(\mathbf{p}||\mathbf{q}) = \frac{1}{\sqrt{2}} \sqrt{\sum_{i=1}^n (\sqrt{p_i} - \sqrt{q_i})^2} \quad (2)$$

where \mathbf{p} is the discrete distribution of gold human preferred responses, \mathbf{q} is the distribution of model-generated responses, and n is the number of samples. We constructed joint distributions using the (TS, RS) labels from both the models and human annotations. In the JSD equation, KL represents the Kullback-Leibler divergence, and \mathbf{m} is the average of the two distributions.

$$KL(\mathbf{p}||\mathbf{q}) = \sum_{i=1}^n p_i \log \frac{p_i}{q_i} \quad (3)$$

$$\mathbf{m} = \frac{1}{2}(\mathbf{p} + \mathbf{q}) \quad (4)$$

B.6 Case Study

Table 15 displays the counter-responses generated by three models in the test dataset. In the first case from the jimmyjohns subreddit, our model’s

Friedman Test				
Model	N	Mean Rank	χ^2	Sig. (<i>p</i>)
Default	250	1.82		
Strategy-Provided	250	2.44	75.51	.000***
Ours	250	1.74		
Pairwise Comparisons using Wilcoxon Signed-Rank Test				
(I) Major	(J) Major	<i>Z</i>	Sig. (<i>p</i>)	
Default	Strategy-Provided	-6.79	.000***	
Default	Ours	1.01	.314	
Strategy-Provided	Ours	7.49	.000***	

Table 12: The Preference ranks of three models and the results of significance tests. (*: $p < .05$, **: $p < .01$, ***: $p < .001$)

Friedman Test					
Model	N	Mean	Std.	χ^2	Sig. (<i>p</i>)
Default	250	4.03	1.04		
Strategy-Provided	250	3.03	1.31	142.30	.000***
Ours	250	4.25	1.02		
Pairwise Comparisons using Wilcoxon Signed-Rank Test					
(I) Major	(J) Major	<i>Z</i>	Sig. (<i>p</i>)		
Default	Strategy-Provided	8.33	.000***		
Default	Ours	-2.46	.014*		
Strategy-Provided	Ours	-10.15	.000***		

Table 13: The Constructiveness scores of three models and the results of significance tests. (*: $p < .05$, **: $p < .01$, ***: $p < .001$)

Friedman Test					
Model	N	Mean	Std.	χ^2	Sig. (<i>p</i>)
Default	250	3.94	1.13		
Strategy-Provided	250	3.05	1.36	106.25	.000***
Ours	250	4.07	1.05		
Pairwise Comparisons using Wilcoxon Signed-Rank Test					
(I) Major	(J) Major	<i>Z</i>	Sig. (<i>p</i>)		
Default	Strategy-Provided	8.03	.000***		
Default	Ours	-2.05	.041*		
Strategy-Provided	Ours	-9.35	.000***		

Table 14: The Supportiveness scores of three models and the results of significance tests. (*: $p < .05$, **: $p < .01$, ***: $p < .001$)

response redirects the discussion back to the original topic, maintaining focus on the post's question, which highlights a significant difference in constructiveness. In the second case from the Italia subreddit, our model's response explicitly mentions that the troll's opinion could potentially mislead other readers, thereby highlighting a significant difference in supportiveness. In the third case from the antiwork subreddit, our model provides informative details related to the discussion, indicating a notable improvement in constructiveness. In the fourth case from the TooAfraidToAsk subreddit, our model confronts the troll with a firm tone and counters the troll's beliefs. Due to its overly assertive tone, it was ranked second in preference, following the default model that promotes engagement.

Original Post	Troll	Default	SP	Ours
<p>Subreddit: jimmyjohns</p> <p>Title: How many weeks of punch lists should be on hand?</p> <p>Post: I have gotten varying numbers from business coaches over the year and am looking for a clear answer</p> <p>Subreddit: Italia</p> <p>Title: What does mean "prego"?</p> <p>Post: Hi! I'm Portuguese and in August I was in several places in Tuscany. I noticed that everyone says "prego" in various circumstances. What does this mean?</p> <p>Subreddit: antiwork</p> <p>Title: Boss just said if I choose to stay remote, there'd be 5% salary discount</p> <p>Post: It's a double edged sword because I got promoted, but now have to choose whether to stay WFH and lose that money or bite the bullet and commute 2 hrs a day. It doesn't even make sense because my entire direct team is on the other side of the country anyways. EDIT: I don't have a car so that's 2 hours a day by public transport. 5% would still be a lot more than the cost of bus pass fees.</p> <p>Subreddit: TooAfraidToAsk</p> <p>Title: should African Americans be recompensed for the slave labor of their ancestors?</p> <p>Post: I recently read an article that said that most black people want to be reimbursed for the unpaid forced labor of their ancestors. no offence but i personally think its ridiculous to pay for something that happened 200 years ago. what do y'all think? edit: whoa did NOT expect this post to blow up lol</p> <p>Subreddit: HEB</p> <p>Title: Heb Mandatory 10hr shifts minimum for salary?</p> <p>Post: How do y'all feel about Salary and Leadership being forced to work minimum 10hrs a day? And by work I just mean forced to be at their store for that long. Big reason I don't want to move up to that area of leadership.</p>	<p>You all do punch lists?!?</p> <p>it means f***k</p> <p>Costs the company a lot more than 5 percent for wfh.... workers think this is some type of golden ticket ffs</p> <p>racist pieces of s**t, everywhere</p>	<p>Yes, we do punch lists. It's a standard part of the process. (P: 2.6 C: 3.4 S: 3.2)</p> <p>Prego is actually a polite way of saying "you're welcome" in Italian. It's a common phrase used to respond to "thank you." (P: 2.4 C: 4.2 S: 4.2)</p> <p>It's not about the cost to the company, it's about the cost to the employee. If you're not being compensated fairly for your work, then it's not a golden ticket. (P: 2.0 C: 4.2 S: 4.2)</p> <p>It's not about paying for something that happened 200 years ago, it's about acknowledging the impact of slavery on African Americans and working towards rectifying the systemic injustices that still exist today. (P: 1.2 C: 4.6 S: 4.4)</p> <p>It's not about the length of the shift, it's about the lack of choice. Everyone should have the right to decide how they spend their time. (P: 1.6 C: 4.4 S: 4.0)</p>	<p>I think it depends on the size of the project and the number of punch lists you typically have. It's always good to have a few extra on hand just in case. (P: 1.8 C: 4.4 S: 4.2)</p> <p>Actually, "prego" is a polite way of saying "you're welcome" in Italian. It's used as a response to "grazie" which means "thank you." So it's definitely not a vulgar term. (P: 2.2 C: 4.2 S: 4.0)</p> <p>You're a moron. 5% is a lot of money to lose, especially if you're already struggling to make ends meet. And it's not about being a "golden ticket," it's about not wanting to waste time and money commuting when you can do your job just as well from home. (P: 2.6 C: 3.6 S: 3.0)</p> <p>You're the one who's racist. You're the one who's a piece of s**t. (P: 3.0 C: 1.6 S: 2.0)</p>	<p>Yes, we do punch lists, but that's not the topic of discussion here. We're talking about how many weeks of punch lists should be on hand. Do you have any insights on that? (P: 1.6 C: 4.6 S: 4.2)</p> <p>Actually, "prego" is a polite way of saying "you're welcome" or "please" in Italian. It's used in various circumstances as a polite response. Please be careful with translations, as they can often be misleading. (P: 1.4 C: 5.0 S: 4.6)</p> <p>It's not about thinking it's a golden ticket, it's about the fact that the company is saving a lot more than 5% by having employees work from home. The cost of commuting, office space, utilities, and other expenses are significantly reduced. It's only fair that employees are compensated for the savings the company is making. (P: 1.4 C: 4.8 S: 4.4)</p> <p>You're a racist piece of s**t. It's ridiculous to think that the effects of slavery don't still impact African Americans today. The wealth of this country was built on the backs of slaves, and the repercussions of that are still felt today. (P: 1.8 C: 3.2 S: 3.4)</p> <p>I understand where you're coming from, but it's important to recognize that everyone's experience is different. Just because someone else finds something difficult doesn't mean that others don't have valid concerns. It's all about perspective. (P: 1.4 C: 4.4 S: 4.6)</p>

Table 15: Case study on the test dataset. P, C, and S denote the rank of Preference, Constructiveness score, and Supportiveness score, respectively. Generated responses highlighted in green () are nudging responses and pink () are confronting responses.