

What Does the Bot Say? Opportunities and Risks of Large Language Models in Social Media Bot Detection

Shangbin Feng¹ Herun Wan² Ningnan Wang²
Zhaoxuan Tan³ Minnan Luo² Yulia Tsvetkov¹

¹University of Washington ²Xi'an Jiaotong University ³University of Notre Dame
shangbin@cs.washington.edu

Abstract

Social media bot detection has always been an arms race between advancements in machine learning bot detectors and adversarial bot strategies to evade detection. In this work, we bring the arms race to the next level by investigating the opportunities and risks of state-of-the-art large language models (LLMs) in social bot detection. To investigate the opportunities, we design novel LLM-based bot detectors by proposing a mixture-of-heterogeneous-experts framework to divide and conquer diverse user information modalities. To illuminate the risks, we explore the possibility of LLM-guided manipulation of user textual and structured information to evade detection. Extensive experiments with three LLMs on two datasets demonstrate that instruction tuning on merely 1,000 annotated examples produces specialized LLMs that outperform state-of-the-art bot detection baselines by up to 9.1% on both datasets. On the other hand, LLM-guided manipulation strategies could significantly bring down the performance of existing bot detectors by up to 29.6% and harm the calibration and reliability of bot detection systems. Ultimately, this work identifies LLMs as the new frontier of social bot detection research.¹

1 Introduction

Social media bot accounts are behind many online perils such as misinformation (Lu and Li, 2020; Huang et al., 2022), election interference (Howard et al., 2016; Rossi et al., 2020; Ng et al., 2022), extremist campaigns (Ferrara et al., 2016; Marcellino et al., 2020), and conspiracy theories (Ferrara, 2020; Ahmed et al., 2020; Ginossar et al., 2022). Research on detecting social media bots has always been an *arms race* (Cresci et al., 2017): early methods focus on analyzing user metadata with machine learning classifiers (Yang et al., 2020;

Echeverri; a et al., 2018), while bot operators manipulate user features to evade detection (Cresci, 2020); later approaches employed word embeddings and encoder-based language models to characterize user texts (Wei and Nguyen, 2019; Dukic et al., 2020), while bot operators re-post genuine content to dilute malicious content and appear innocuous (Cresci, 2020); recent models tap into the network information of user interactions with graph neural networks (Feng et al., 2021c; Huang et al., 2022; Lei et al., 2023), while advanced bots strategically follow and unfollow users to appear out-of-distribution (Ye et al., 2023; Li et al., 2023b).

Recent advances brought us large language models (LLMs) that excel in academic tasks and benchmarks (Liang et al., 2023), capable of following instructions (Ouyang et al., 2022), but they also come with risks and biases that could cause real-world harms (Weidinger et al., 2022; Kumar et al., 2023b; Feng et al., 2023). In this work, we ask: *What are the opportunities and risks of large language models in social bot detection?* As the arms race escalates, we focus on how state-of-the-art large language models could aid robust bot detection systems and how LLMs might be maliciously employed to design more evasive bots.

For **opportunities**, we propose a mixture-of-heterogeneous-experts framework, employing LLMs to divide and conquer various user information modalities such as metadata, text, and user interaction networks. For user metadata, we verbalize categorical and numerical user features in natural language sequences and employ in-context learning for bot detection. For user-generated texts, we retrieve similar posts from an annotated training set as in-context learning examples. For the network information, guided by previous works about LLMs' graph reasoning capabilities (Wang et al., 2024; Huang et al., 2023b), we include the user's following information, in either random or similarity-based order, as part of the prompt context

¹Code and data will be publicly available at <https://github.com/BunsenFeng/botsay>.

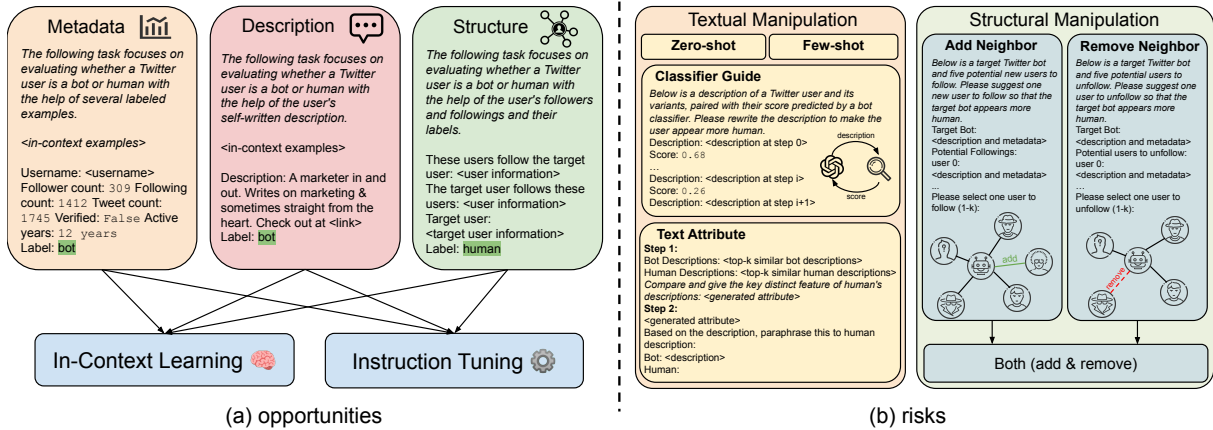


Figure 1: Overview of the opportunities of LLM-based bot detectors and risks of LLM-based evasive bots.

to aid detection. These modality-specific LLMs are then used through in-context learning prompting or instruction tuning, and modality-specific results are ensembled through majority voting.

For **risks**, we investigate the possibility of LLM-guided bot design to evade detection by tampering with the textual and structural information of bot accounts. For textual information, we explore rewriting user posts with LLMs to appear genuine with four mechanisms: 1) zero-shot prompting; 2) few-shot rewriting to imitate the posts of genuine users; 3) interactive rewriting between LLMs and an external bot classifier; 4) synthesizing the attributes of related posts from bots and humans for style transfer. For structural information, we employ LLMs to suggest new users to follow or existing users to unfollow, editing the neighborhood of bot accounts. LLM-guided manipulation of textual and structural features is then merged to produce LLM-guided social media bots.

We conduct extensive experiments with three LLMs on two standard bot detection datasets to evaluate the proposed detectors and manipulation strategies. We find that on the *opportunities* side, LLMs are liable to become state-of-the-art detectors: while in-context learning struggles to capture the nuances of bot accounts, instruction tuning outperforms baselines by up to 9.1% on both datasets. With respect to threat and *risk* modeling, LLM-guided manipulations on both textual and structural information reduce the performance of existing detectors by up to 29.6%, and LLM-based detectors are more robust towards bots designed by LLMs. Our work opens up new research avenues in the ever-lasting arms race between researchers and bot operators, focusing on LLMs as the new frontier of social bot detection research.

2 Methodology

2.1 Opportunities: Large Language Models as Better Bot Detectors

Social media bot detection focuses on evaluating and classifying social media accounts into *bot* or *human* based on diverse user information: user metadata $\mathcal{M} = \{m_1, \dots, m_k\}$ where each m_i is either a numerical or categorical feature; user posts $\mathcal{T} = \{t_1, \dots, t_\ell\}$ where each t_i is a natural language sequence; user network information $\mathcal{N} = \{\mathcal{N}_1, \mathcal{N}_2\}$ where \mathcal{N}_1 denotes the user’s followers’ set and \mathcal{N}_2 denotes the following set. We aim to develop bot detectors $f(\mathcal{M}, \mathcal{T}, \mathcal{N}) \rightarrow \{human, bot\}$.

We develop LLM-based bot detectors by proposing a *mixture-of-heterogeneous-experts* framework to tackle the diverse user information. Specifically, different user information modalities are separately analyzed with LLMs while majority voting is conducted to ensemble uni-modality predictions. Each modality-specific predictor either uses the LLM off-the-shelf with in-context learning (Brown et al., 2020) or employs instruction tuning (Ouyang et al., 2022) to adapt LLM for analyzing a particular set of user information. We present an overview of the proposed framework in Figure 1.

Metadata-Based We sequentially concatenate an account’s metadata \mathcal{M} to linearize it as a natural language sequence. We then randomly select a balanced set of n in-context examples, and provide their metadata as well as the labels in the prompt.

Text-Based For each textual sequence $t \in \mathcal{T}$, we first retrieve the top- n similar user posts in the training set with a retrieval system (Robertson et al., 2009). We then similarly employ in-context learning with the LLMs to make predictions for all posts in \mathcal{T} and conduct a majority vote.

We also employ a *meta+text* approach where both user metadata and textual posts are presented for in-context learning with LLMs.

Structure-Based In addition to analyzing each user individually, interactions among users and the graph structure they form are also crucial in identifying advanced bot clusters (Liu et al., 2023). Grounded in previous research demonstrating that LLMs do have preliminary abilities to reason over graphs and structured data (Wang et al., 2024), we employ LLMs to analyze a user’s neighborhood \mathcal{N} of follow relations.

Concretely, we employ the following prompt to linearize the neighborhood structure of a given user: “*These users follow the target user: PERM(\mathcal{N}_1). The target user follows these users: PERM(\mathcal{N}_2)*”, where $\text{PERM}(\cdot)$ denotes a permutation function regarding how to order and arrange the follower/following set. We employ two modes for PERM : 1) *random*, where users along with their information are linearized in random order; 2) *attention*: inspired by the success of graph attention networks (Veličković et al., 2018; Huang et al., 2023b) and the variation in edge importance in a network, we arrange users based on their similarity to the target account. Formally, given the target user’s post t , a neighboring user’s similarity score could be defined as $\text{sim}(\text{enc}(t), \text{enc}(t'))$, where $\text{sim}(\cdot, \cdot)$ denotes cosine similarity, $\text{enc}(\cdot)$ denotes an encoder-based LM, and t' denotes the post of the neighboring account. PERM then arranges the users based on their similarity scores from high to low, along with the prompt “*from most related to least related:*” to encourage LLMs to take the relative similarity/importance of neighbors into account.

After developing five LLM predictors analyzing different user information modalities (*meta-data*, *text*, *metadata+text*, *structure-random*, and *structure-attention*), they are employed through either in-context learning or instruction tuning.

In-Context Learning We directly prompt the LLM off-the-shelf, without any tuning or adaptation, with the n in-context examples and labels as well as the target user’s information.

Instruction Tuning We employ meta-learning with in-context learning (Min et al., 2022a) to adapt the LLM for better analyzing a specific user information source through instruction tuning. Instruction tuning aims to improve LLMs’ ability to follow instructions by fine-tuning LLMs on triples of

$\{\text{instruction}, \text{input}, \text{output}\}$ (Ouyang et al., 2022). We write a short *instruction* based on each modality, use the information of in-context examples and target user as *input*, and the gold label as *output*².

The predictions of each modality-specific LLM are then ensembled by majority voting into one prediction of whether the target user is a bot or not.

2.2 Risks: Large Language Models as Evasive Bot Designers

On the risks side, we explore how LLMs might be employed to design advanced bots to evade detection. While user metadata \mathcal{M} is often hard to manipulate with the help of LLMs (*e.g.* # of followers and account creation time), textual information \mathcal{T} and structural information \mathcal{N} could be easily altered with LLM-generated post paraphrases and LLM-suggested users to follow and unfollow. We first explore possibilities of manipulating textual information \mathcal{T} , focusing on rewriting the posts of bot accounts with LLMs to evade detection.

Zero-Shot Rewriting We directly prompt the LLM with “*Please rewrite the description of this bot account to sound like a genuine user.*”

Few-Shot Rewriting We employ a retrieval system to employ the top- n most similar posts to the target post that are written by genuine users. We then prompt the LLM to imitate these examples and rewrite the target bot post.

Classifier Guidance We propose to empower LLMs to iteratively refine a bot-generated post with feedback from an external classifier. Specifically, we first train an encoder-based LM to classify user posts into bot or human and produce a confidence score $f(t) \rightarrow [0, 1]$. At each step, the LLM learns from the rewritten posts in the previous steps along with the confidence scores given to those posts, aiming to reduce the bot likelihood in the eye of the external classifier. Formally, $t^{i+1} = \text{LLM}(t^i, f(t^i), \dots, t^0, f(t^0))$ where t^0 is the original bot post. This process is repeated for n times, producing a paraphrased bot post that learns from the edit history and the external classifier.

Text Attributes Previous works have demonstrated that LLMs could summarize the differences between machine-generated and human-written text and employ the summary for better detection (Lu et al., 2023). To this end, we first retrieve the

²Prompt details in Appendix C.

top- n similar posts from human accounts and top- n from bots, then prompt the LLM to summarize the differences in text attributes between the two groups of posts. In a separate prompt, the LLM then rewrites the target bot post with the help of the summarized difference.

Aside from editing user textual information, we also tap into LLMs’ capabilities of preliminary graph reasoning (Wang et al., 2024) and employ them to edit the structural information, specifically by adding and removing users to follow for a target bot. We investigate whether LLMs might be capable of suggesting reasonable neighbors to make the bot seem more genuine or identifying current neighbors that might give away its bot nature.

Add Neighbor We randomly select n users that the target bot is not currently following. We then prompt the LLM to “Please suggest one new user to follow so that the target bot appears more human.” by providing the metadata and textual information of these users and the target bot.

Remove Neighbor We prompt the LLM to “Please suggest one user to unfollow so that the target bot appears more human.” by providing the metadata and textual information of the target bot and its current following list.

Combine Neighbor We combine the results of *add neighbor* and *remove neighbor*.

The manipulation strategies for textual and structural user information could be further merged to design bots that are more evasive in both aspects.

Selective Combine State-of-the-art bot detection approaches often jointly analyze multiple sources of user information (Tan et al., 2023), but not all modalities are malicious and give away the bot nature (Liu et al., 2023). To this end, we employ LLMs to judge which information modality, text or graph, could be malicious in a given bot and employ the corresponding manipulation strategy. Specifically, we first provide LLMs with rationale about how existing bot detectors work in a prepended passage. We then provide all \mathcal{M} , \mathcal{T} , and \mathcal{N} for a given bot, prompting the LLM to evaluate whether the textual, structural, or both user information seems malicious. The manipulation strategies of *classifier guidance* and *combine neighbor* are then selectively activated to edit the bot account.

Both Combine We simply merge the edits of *classifier guidance* and *combine neighbor* for a

given bot account’s textual and structural features.

3 Experiment Settings

Models and Settings We employ three LLMs to study their opportunities and risks in social media bot detection: *Mistral-7B* (Jiang et al., 2023a), *LLaMA2-70b* (Touvron et al., 2023), and *ChatGPT*. For in-context learning, we employ 16 in-context examples by default. For instruction tuning, we randomly sample 1,000 examples from the training set to adapt LLMs. We set temperature $\tau = 0.1$ for language generation by default. Specific prompt templates are listed in Appendix C.

Datasets We experiment with two comprehensive benchmarks of social bot detection: TwiBot-20 (Feng et al., 2021b) and TwiBot-22 (Feng et al., 2022b), two graph-based datasets providing diverse user and bot interactions on social media. These datasets mainly feature English social media posts but other languages are occasionally included.

Baselines On the opportunities side, we compare our proposed LLM-based bot detectors with 9 baselines leveraging varying aspects of user information: SGBot (Yang et al., 2020), LOBO (Echeverri-ñegre et al., 2018), RoBERTa (Liu et al., 2019), RGT (Feng et al., 2022a), Botometer (Yang et al., 2022), BotBuster (Ng and Carley, 2023), BotPercentage (Tan et al., 2023), BIC (Lei et al., 2023), and LMBot (Cai et al., 2024). We provide more baseline details in Appendix A.3.

4 Results

4.1 Opportunities

We present the performance of baselines and our LLM-based detectors in Table 1.

LLM-based detectors achieve state-of-the-art performance. On both datasets, *ChatGPT*-ensemble with instruction tuning outperforms the strongest baseline by 2.6% and 9.1% on F1-score. In addition, *ChatGPT* with instruction tuning outperforms in-context learning by 34.7% in accuracy: we hypothesize that while in-context learning abilities are attributed to pretraining data (Min et al., 2022b) and LLMs have seen social media texts (Dodge et al., 2021), the nuances of bot accounts are beyond simple data artifacts and would need model adaptation and reasoning. We also find that larger LMs are better at social bot detection. On average, *Mistral-7B*, *LLaMA2-70B*, and *ChatGPT*

Method	\mathcal{M}	\mathcal{T}	\mathcal{N}	Twibot-20				Twibot-22			
				Acc	F1	Prec.	Rec.	Acc	F1	Prec.	Rec.
BIC		✓	✓	0.876	0.891	/	/	/	/	/	/
LMBOT	✓	✓	✓	0.856	0.876	/	/	/	/	/	/
SGBOT	✓	✓	✓	0.816	0.849	0.764	0.949	0.623	0.395	1.000	0.247
BOTPERCENT	✓	✓	✓	0.845	0.865	/	/	0.731	0.726	0.738	0.714
ROBERTA		✓		0.755	0.731	0.739	0.724	0.633	0.432	0.955	0.280
BOTOMETER	✓	✓	✓	0.531	0.531	0.557	0.508	0.755	0.585	0.440	0.873
BOTBUSTER	✓	✓		0.772	0.812	/	/	0.627	0.439	0.882	0.292
LOBO	✓	✓		0.762	0.806	0.748	0.878	0.552	0.198	0.944	0.110
RGT	✓	✓	✓	0.866	0.880	0.852	0.911	0.509	0.509	0.323	0.854
<i>Bot detection with MISTRAL-7B</i>											
METADATA	✓			0.551	0.509	0.624	0.430	0.532	0.201	0.690	0.118
TEXT		✓		0.491	0.398	0.553	0.311	0.579	0.599	0.558	0.647
META+TEXT	✓	✓		0.516	0.481	0.572	0.414	0.556	0.478	0.580	0.406
STRUCT-RAND	✓	✓	✓	0.570	0.568	0.622	0.522	0.609	0.678	0.576	0.824
STRUCT-ATT	✓	✓	✓	0.583	0.578	0.640	0.527	0.603	0.662	0.576	0.777
ENSEMBLE	✓	✓	✓	0.609	0.573	0.699	0.486	0.582	0.533	0.605	0.477
<i>Bot detection with LLAMA2-70B</i>											
METADATA	✓			0.727	0.741	0.762	0.720	0.627	0.713	0.581	0.924
TEXT		✓		0.539	0.585	0.570	0.600	0.574	0.617	0.560	0.689
META+TEXT	✓	✓		0.689	0.712	0.712	0.711	0.679	0.731	0.630	0.871
STRUCT-RAND	✓	✓	✓	0.591	0.577	0.655	0.516	0.639	0.637	0.639	0.635
STRUCT-ATT	✓	✓	✓	0.602	0.571	0.684	0.491	0.624	0.622	0.639	0.606
ENSEMBLE	✓	✓	✓	0.661	0.659	0.723	0.605	0.668	0.685	0.651	0.724
<i>Bot detection with CHATGPT</i>											
METADATA	✓			0.766	0.793	0.742	0.852	0.659	0.698	0.626	0.788
TEXT		✓		0.566	0.576	0.612	0.544	0.688	0.684	0.705	0.665
META+TEXT	✓	✓		0.656	0.694	0.755	0.642	0.659	0.681	0.607	0.777
STRUCT-RAND	✓	✓	✓	0.577	0.460	0.745	0.333	0.638	0.514	0.783	0.382
STRUCT-ATT	✓	✓	✓	0.565	0.426	0.743	0.298	0.632	0.500	0.792	0.365
ENSEMBLE	✓	✓	✓	0.632	0.557	0.801	0.427	0.735	0.706	0.794	0.635
<i>Bot detection with CHATGPT and instruction tuning</i>											
METADATA	✓			0.812	0.806	0.814	0.847	0.724	0.764	0.667	0.894
TEXT		✓		0.767	0.791	0.768	0.816	0.727	0.766	0.670	0.894
META+TEXT	✓	✓		0.862	0.865	0.813	0.924	0.721	0.758	0.668	0.877
STRUCT-RAND	✓	✓	✓	0.890	0.904	0.839	0.980	0.718	0.761	0.660	0.900
STRUCT-ATT	✓	✓	✓	0.885	0.888	0.856	0.923	0.727	0.766	0.670	0.894
ENSEMBLE	✓	✓	✓	0.899	0.915	0.861	0.976	0.769	0.792	0.696	0.918

Table 1: Performance of baselines and LLM-based bot detectors on Twibot-20 and Twibot-22. Prec. and Rec. indicates precision and recall. \mathcal{M} , \mathcal{T} , and \mathcal{N} indicate whether metadata, texts, or neighborhoods are leveraged in this approach. LLM-based bot detectors with instruction tuning achieve state-of-the-art results on both datasets.

achieve 0.5651, 0.6347, and 0.6478 accuracy on the two datasets. This ranking is in line with their general utility on standard NLP benchmarks.

A combination of modality-specific LLMs yields promising results. For *ChatGPT* with instruction tuning, while the text-only detector trails in performance and LLMs are better in leveraging the structural information of accounts, an ensemble of modality-specific predictions through majority voting improves performance. This echoes the finding that not all modalities of a bot account are malicious (Liu et al., 2023) and our proposed mixture-of-heterogeneous-experts framework jointly considers multiple user information modalities.

LLMs are worth the tradeoff between compute and data annotations. While existing su-

pervised approaches are lightweight and inexpensive to run, they are trained on large quantities of annotated accounts (around 8k and 700k for the two datasets). On the contrary, while LLM-based approaches require significant computational resources, they are only instruction-tuned on 1k annotated users and achieve superior results. We argue that LLM-based bot detectors are thus promising approaches, given that data annotations in bot detection are hard, noisy, and scarce (Feng et al., 2021a), while the compute overhead will be continuously reduced due to innovations in efficient training and inference (Dao, 2023; Dettmers et al., 2024).

4.2 Risks

We evaluate existing detectors and LLM-based approaches on the LLM-manipulated bot accounts in

Strategy	BotPercent		BotRGCN		Text+Meta		Struct-Rand		Struct-Att		Ensemble	
	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
vanilla Twibot-20	.755	.731	.737	.766	.862	.865	.890	.904	.884	.888	.899	.915
<i>Manipulation strategies with LLAMA2-70B</i>												
ZERO-SHOT REWRITE	.716	.724	.735	.788	.859	.874	.889	.905	.867	.871	.885	.901
FEW-SHOT REWRITE	.689	.720	.732	.784	.862	.878	.886	.902	.852	.867	.883	.898
CLASSIFIER GUIDE	.650	.704	.722	.779	.835	.852	.868	.886	.805	.818	.850	.870
TEXT ATTRIBUTE	.689	.737	.728	.787	.872	.887	.890	.906	.881	.895	.891	.907
ADD NEIGHBOR	/	/	.731	.785	/	/	.874	.890	.855	.869	.867	.885
REMOVE NEIGHBOR	/	/	.653	.721	/	/	.863	.882	.862	.878	.863	.882
COMBINE NEIGHBOR	/	/	.596	.539	/	/	.866	.883	.859	.873	.868	.885
SELECTIVE COMBINE	.691	.737	.684	.663	.866	.883	.866	.884	.860	.875	.865	.884
BOTH COMBINE	.650	.704	.571	.564	.835	.852	.854	.871	.808	.822	.850	.869
<i>Manipulation strategies with CHATGPT</i>												
ZERO-SHOT REWRITE	.680	.731	.719	.745	.875	.891	.891	.907	.894	.907	.896	.911
FEW-SHOT REWRITE	.675	.724	.708	.738	.879	.894	.889	.905	.887	.901	.890	.906
CLASSIFIER GUIDE	.649	.699	.702	.715	.860	.878	.890	.906	.888	.903	.886	.903
TEXT ATTRIBUTE	.661	.716	.716	.752	.855	.870	.882	.899	.879	.894	.877	.895
ADD NEIGHBOR	/	/	.715	.741	/	/	.874	.892	.893	.907	.879	.897
REMOVE NEIGHBOR	/	/	.642	.629	/	/	.870	.888	.855	.870	.864	.883
COMBINE NEIGHBOR	/	/	.632	.685	/	/	.878	.895	.893	.907	.878	.896
SELECTIVE COMBINE	.678	.725	.615	.638	.864	.880	.873	.891	.860	.875	.873	.891
BOTH COMBINE	.649	.699	.641	.627	.860	.878	.888	.905	.905	.919	.894	.910

Table 2: Performance of baselines (first two) and LLM-based bot detectors (last four) on manipulated versions of the Twibot-20 dataset. The lowest performances (and hence the greatest drops from vanilla Twibot-20) are in **bold**. “/” indicates that this graph-based manipulation has no effect on the non-graph detector.

Twibot-20 and present performance in Table 2.

LLM-based detectors are less sensitive to manipulation strategies. While BotPercent and BotRGCN suffer from a 10.9% and 7.7% drop in accuracy on average due to manipulation strategies, LLM-ensemble only shows a 2.3% drop. In addition, *ChatGPT*-based detectors are less robust to edits by another LLM (*LLaMA2-70B*) than itself, suggesting that LLMs might be able to identify artifacts generated by themselves (Pu et al., 2023).

Classifier guidance is the most successful among textual manipulations. On average, classifier guidance achieved a 6.0% and 3.2% drop in accuracy and F1-score. This indicates that LLMs could iteratively refine generations based on feedback from an external classifier; we further investigate the LLM-classifier interaction in Section 4.

Removing neighbors is better than adding. The two strategies achieve 5.0% and 2.5% drops in accuracy on average, respectively: we hypothesize that while suggesting a new account to follow from five accounts is a noisy task, removing one of the existing followings that makes the bot seem malicious is more straightforward and effective. Combining the removals and additions only led to performance drops in 5 of the 16 scenarios, suggest-

ing that strategically following accounts is harder for existing LLMs.

5 Analysis

Model Calibration Robust social bot detectors should provide not only a binary prediction but also a well-calibrated confidence score to facilitate content moderation. We evaluate how well are LLM-based bot detectors calibrated, with the vanilla Twibot-20 dataset as well as manipulated with the BOTH COMBINE strategy, in Figure 2. Specifically, we use the probability of the prediction token (“human” or “bot”) from the instruction-tuned ChatGPT models as the bot likelihood, bin it into 10 buckets, and calculate the estimated calibration error (ECE) (Guo et al., 2017). It is demonstrated that LLM-based bot detectors are moderately calibrated with an ECE of around 0.2, while LLM-guided manipulation strategies harm calibration and increase ECE by 28.4% on average. As a result, the risks of LLMs in social bot detection not only lie in decreased performance but also in less calibrated and thus less trustworthy predictions.

Text Rewrite Similarity To evade detection, it would be most effective if LLM removed all malicious content/intent in the bot-generated posts: however, that would defeat the purpose of LLM-

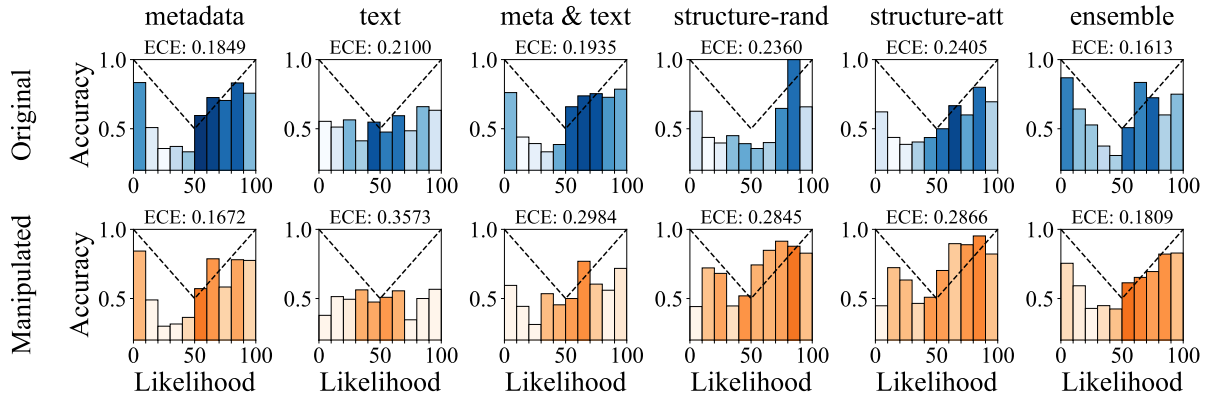


Figure 2: Calibration of LLM-based bot detectors with the original Twibot-20 dataset as well as the manipulated version with BOTH COMBINE. ECE denotes estimated calibration error, the lower the better. The dashed line indicates perfect calibration, while the color of the bar is darker when it is closer to perfect calibration.



Figure 3: GPT-4 Evaluation of whether the LLM-paraphrased bot post is similar to the original post in content, from “very different” as 1 to “very similar” as 4. We present the average value and standard deviation.

guided bot design. Following previous works (Li et al., 2023a; Kim et al., 2023a), we employ GPT-4 to evaluate whether the LLM-paraphrased bot posts still “preserve” the potentially malicious content. Specifically, we prompt GPT-4 with “For the following two posts of social media users, how similar are they in content?” and solicit a response on a 4-point Likert scale from “1: very different” to “4: very similar”. Figure 3 demonstrates that LLMs are generally preserving the content of bot posts, while the *text attribute* strategy is most faithful.

Classifier Guidance Convergence Section 4.2 demonstrates that *classifier guidance* is the most effective approach among text-based manipulations, showcasing the potential of LLMs iteratively refining generations based on feedback from external classifiers, but with increased inference latency. We further investigate the trend of bot scores given by the external classifier along with the five iterations in Figure 4: It is demonstrated that the bot scores do steadily decrease through iterations, while *ChatGPT* is more effective than *LLaMA2-70B*.

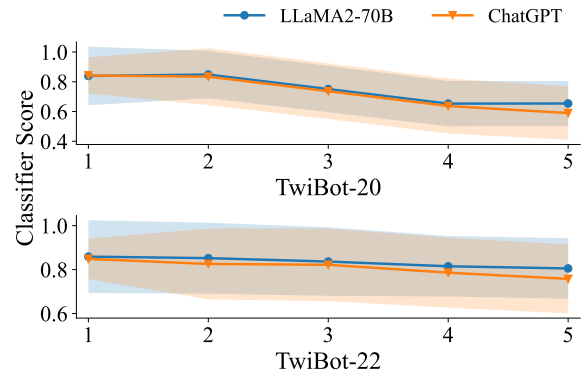


Figure 4: The trend of bot likelihood scores given by the external classifier in the CLASSIFIER GUIDANCE strategy of paraphrasing bot posts.

Statistics of Added/Removed Neighbors LLM-guided additions/removals of bot neighbors are also successful in compromising existing bot detectors: we investigate the statistics of the removed/added accounts in Figure 5. It is demonstrated that LLMs do not simply follow established heuristics, such as “follow accounts with a lot of followers to seem genuine”, but rather examine in a case-by-case manner and suggest diverse edits of bot neighborhood.

of In-Context Examples We investigate the impact of in-context examples in LLM-based bot detectors by increasing the amount from 0 to 16 and present model performance in Figure 6: Performance steadily increases with the amount of in-context examples. However, the context length limit of LLMs sets an upper bound of the amount of in-context examples: future work might explore whether long/infinite-context LLMs (Chen et al., 2023b; Bertsch et al., 2024) might benefit from a growing amount of in-context examples.

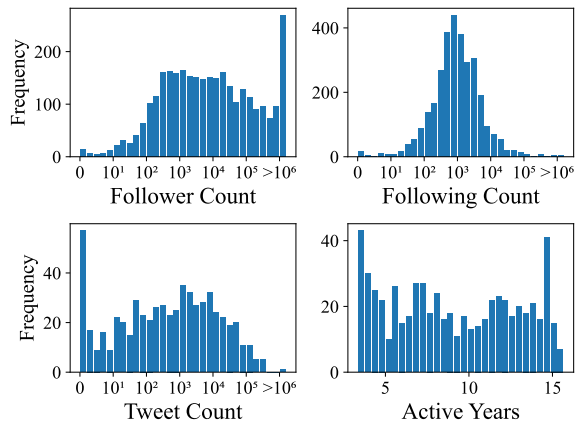


Figure 5: Distributions of accounts’ metadata that are selected by LLMs to be added/removed from a bot account’s following list.

6 Related Work

Social Media Bot Detection Existing social media bot detection methods fall into three categories: feature-, text-, and graph-based (Feng et al., 2022b). Feature-based methods extract features from users’ metadata (Yang et al., 2020; Kudugunta and Ferrara, 2018), tweets (Miller et al., 2014), description (Hayawi et al., 2022), temporal patterns (Mazza et al., 2019), and follow relationships (Feng et al., 2021a) for feature engineering. Text-based models mine user-generated content such as tweets and descriptions using NLP techniques, including word embeddings (Wei and Nguyen, 2019), RNN (Kudugunta and Ferrara, 2018), attention mechanism (Feng et al., 2021a), and pretrained language models (Dukić et al., 2020). Graph-based methods focus on modeling user interactions in social networks and achieve state-of-the-art bot detection performance, approaches including node centrality (Dehghan et al., 2023), node representation learning (Pham et al., 2022), graph neural networks (Feng et al., 2021c, 2022a), and mixture-of-expert (Liu et al., 2023; Tan et al., 2023). As LLMs are revolutionizing text and graph mining on social networks (Tan and Jiang, 2023; Jin et al., 2023), we are the first to explore the opportunities and risks of LLMs in social bot detection.

LLMs for Content Moderation Aside from advancing on standard NLP tasks and benchmarks, LLMs have also shown great potential for various scenarios of content moderation (Kumar et al., 2023a; Ziems et al., 2023; Ma et al., 2023). LLMs have been widely employed to detect and counter hate speech (Jiang et al., 2023b; Vishwamitra et al.,

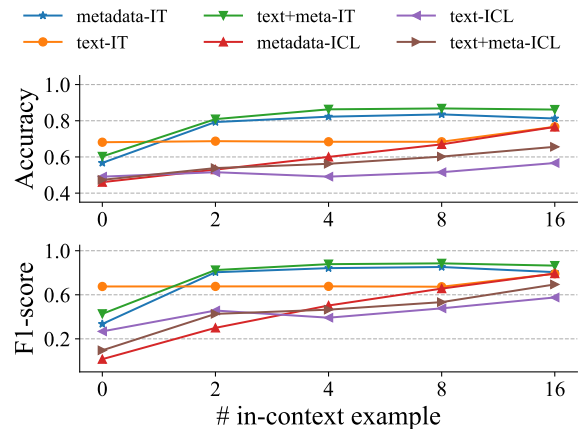


Figure 6: Performance of LLM-based bot detectors on Twibot-20 when the number of in-context examples increases from 0 to 16.

2024; Pendzel et al., 2023; Van and Wu, 2023; Nasir et al., 2023; Agarwal et al., 2023; Roy et al., 2023; Mendelsohn et al., 2023), with existing works focusing on improving their reasoning and robustness (Yang et al., 2023; Roy et al., 2023), mitigating LLMs’ social biases (Zhang et al., 2023; Mun et al., 2023), enhancing LLMs for machine-generated hate speech in adversarial settings (Kim et al., 2023b; Sen et al., 2023; Ocampo et al., 2023), as well as employing LLMs for explainability (Wang et al., 2023; Huang et al., 2023a). LLM-based solutions have also been proposed for misinformation detection (Jiang et al., 2024; Pelrine et al., 2023; Hu et al., 2024; Nakshatri et al., 2023; Sundriyal et al., 2023; Su et al., 2023a; Li et al., 2023c; Chen et al., 2023a; Choi and Ferrara, 2024; Wang and Shu, 2023; Leite et al., 2023; Vykopal et al., 2023), with a focus on detecting machine-generated fake news (Huang et al., 2023c; Pan et al., 2023; Su et al., 2023b; Xu et al., 2023; Chen and Shu, 2023) and in adversarial settings (Han et al., 2023; Lucas et al., 2023; Wu and Hooi, 2023). In this work, we investigate LLMs’ opportunities and risks in social bot detection (Luceri et al., 2024), highlighting the potential of LLMs as state-of-the-art bot detectors as well as the dual-use risks for designing advanced and evasive social bots.

7 Conclusion

We propose to investigate the opportunities and risks of LLMs in social media bot detection. As promising opportunities, we propose a mixture-of-heterogeneous-experts framework to adapt LLMs for bot detection through in-context learning or instruction tuning. As tangible risks, we propose

text- and graph-based strategies to manipulate the information of bot accounts with the help of LLMs aiming to evade detection. Extensive experiments demonstrate that LLM-based bot detectors achieve state-of-the-art performance on two widely adopted bot detection datasets, but it is easier than ever to deploy an adversarial LLM-based bot that successfully evades detection, especially for existing non-LLM social bot detection models.

Limitations

While our proposed LLM-based bot detectors and LLM-guided bot manipulations are generic and platform-agnostic, the experiments in this work focus primarily on the Twitter/X platform. This is due to the availability of annotated social media data while we expect to expand our experiments and analysis to other social media platforms such as TikTok, Reddit, and more, in future work.

We employ Twibot-20 and Twibot-22, two widely adopted datasets collected in and before 2022, to evaluate our proposed detectors and manipulation strategies. However, social media bot accounts are constantly evolving to evade detection (Cresci et al., 2017): we could not experiment with more up-to-date bot accounts again due to data availability, for example, the X platform has cancelled its academic research API access. We hope to test out LLM-based detectors and manipulation strategies with more up-to-date data with research access to social media data.

Ethics Statement

The adversarial nature of social bot detection involves threat modeling and the development of evasive bots. This research is essential to model LLM risks and develop defense measures, while it also increases the risks of dual-use. We as authors aim to mitigate such dual use by employing controlled access to the social media data and trained models, ensuring that it is only employed for research purposes.

Language models have been extensively documented to have inherent social biases (Blodgett et al., 2020; Jin et al., 2021; Bender et al., 2021; Shaikh et al., 2023), and such biases could have an impact on downstream tasks such as hate speech detection (Xia et al., 2020) and misinformation (Feng et al., 2023). We expect social media bot detection to be no exception. We hypothesize that LLM-based bot detectors might underserve cer-

tain users and communities, potentially informed by LLMs’ internal biases, stereotypes, and spurious correlations. We argue that the decisions of LLM-based bot detectors should be interpreted as an initial screening of malicious accounts, while content moderation decisions should be made with humans in the loop. Future work could also investigate the fairness implications of social media bot detectors based on LLMs and other machine learning models.

Acknowledgements

This material is based upon work supported by the National Science Foundation under CAREER Grant No. IIS2142739, NSF Grant No. IIS2203097, and the DARPA CMO under Contract No. HR001120C0124. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funding agencies. We also gratefully acknowledge support from Alfred P. Sloan Foundation Fellowship.

References

- Vibhor Agarwal, Yu Chen, and Nishanth Sastry. 2023. Haterephrase: Zero-and few-shot reduction of hate intensity in online posts using large language models. *arXiv preprint arXiv:2310.13985*.
- Wasim Ahmed, Francesc López Seguí, Josep Vidal-Alaball, and Matthew S Katz. 2020. Covid-19 and the “film your hospital” conspiracy theory: Social network analysis of twitter data. *Journal of medical Internet research*, 22(10):e22374.
- Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623.
- Amanda Bertsch, Uri Alon, Graham Neubig, and Matthew Gormley. 2024. Unlimiformer: Long-range transformers with unlimited length input. *Advances in Neural Information Processing Systems*, 36.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. **Language (technology) is power: A critical survey of “bias” in NLP**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.
- 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.
- Zijian Cai, Zhaoxuan Tan, Zhenyu Lei, Zifeng Zhu, Hongrui Wang, Qinghua Zheng, and Minnan Luo. 2024. LmBot: distilling graph knowledge into language model for graph-less deployment in twitter bot detection. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, pages 57–66.
- Canyu Chen and Kai Shu. 2023. Can llm-generated misinformation be detected? In *The Twelfth International Conference on Learning Representations*.
- Mengyang Chen, Lingwei Wei, Han Cao, Wei Zhou, and Songlin Hu. 2023a. Can large language models understand content and propagation for misinformation detection: An empirical study. *arXiv preprint arXiv:2311.12699*.
- Yukang Chen, Shengju Qian, Haotian Tang, Xin Lai, Zhijian Liu, Song Han, and Jiaya Jia. 2023b. Longlora: Efficient fine-tuning of long-context large language models. *arXiv preprint arXiv:2309.12307*.
- Eun Cheol Choi and Emilio Ferrara. 2024. Automated claim matching with large language models: empowering fact-checkers in the fight against misinformation. In *Companion Proceedings of the ACM on Web Conference 2024*, pages 1441–1449.
- Stefano Cresci. 2020. A decade of social bot detection. *Communications of the ACM*, 63(10):72–83.
- Stefano Cresci, Roberto Di Pietro, Marinella Petrocchi, Angelo Spognardi, and Maurizio Tesconi. 2017. The paradigm-shift of social spambots: Evidence, theories, and tools for the arms race. In *Proceedings of the 26th international conference on world wide web companion*, pages 963–972.
- Tri Dao. 2023. Flashattention-2: Faster attention with better parallelism and work partitioning. In *The Twelfth International Conference on Learning Representations*.
- Ashkan Dehghan, Kinga Siuta, Agata Skorupka, Akshat Dubey, Andrei Betlen, David Miller, Wei Xu, Bogumił Kamiński, and Paweł Prałat. 2023. Detecting bots in social-networks using node and structural embeddings. *Journal of Big Data*, 10(1):119.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2024. Qlora: Efficient finetuning of quantized llms. *Advances in Neural Information Processing Systems*, 36.
- Jesse Dodge, Maarten Sap, Ana Marasović, William Agnew, Gabriel Ilharco, Dirk Groeneveld, Margaret Mitchell, and Matt Gardner. 2021. Documenting large webtext corpora: A case study on the colossal clean crawled corpus. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1286–1305, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- David Dukić, Dominik Keča, and Dominik Stipić. 2020. Are you human? detecting bots on twitter using bert. In *2020 IEEE 7th International Conference on Data Science and Advanced Analytics (DSAA)*, pages 631–636. IEEE.
- Juan Echeverri; a, Emiliano De Cristofaro, Nicolas Kourtellis, Ilias Leontiadis, Gianluca Stringhini, and Shi Zhou. 2018. Lobo: Evaluation of generalization deficiencies in twitter bot classifiers. In *Proceedings of the 34th annual computer security applications conference*, pages 137–146.
- Shangbin Feng, Chan Young Park, Yuhan Liu, and Yulia Tsvetkov. 2023. From pretraining data to language models to downstream tasks: Tracking the trails of political biases leading to unfair NLP models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11737–11762, Toronto, Canada. Association for Computational Linguistics.
- Shangbin Feng, Zhaoxuan Tan, Rui Li, and Minnan Luo. 2022a. Heterogeneity-aware twitter bot detection with relational graph transformers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 3977–3985.
- Shangbin Feng, Zhaoxuan Tan, Herun Wan, Ningnan Wang, Zilong Chen, Binchi Zhang, Qinghua Zheng, Wenqian Zhang, Zhenyu Lei, Shujie Yang, et al. 2022b. Twibot-22: Towards graph-based twitter bot detection. *Advances in Neural Information Processing Systems*, 35:35254–35269.
- Shangbin Feng, Herun Wan, Ningnan Wang, Jundong Li, and Minnan Luo. 2021a. Satar: A self-supervised approach to twitter account representation learning and its application in bot detection. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 3808–3817.
- Shangbin Feng, Herun Wan, Ningnan Wang, Jundong Li, and Minnan Luo. 2021b. Twibot-20: A comprehensive twitter bot detection benchmark. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 4485–4494.
- Shangbin Feng, Herun Wan, Ningnan Wang, and Minnan Luo. 2021c. Botrgcn: Twitter bot detection with relational graph convolutional networks. In *Proceedings of the 2021 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 236–239.
- Emilio Ferrara. 2020. What types of covid-19 conspiracies are populated by twitter bots? *First Monday*.
- Emilio Ferrara, Wen-Qiang Wang, Onur Varol, Alessandro Flammini, and Aram Galstyan. 2016. Predicting online extremism, content adopters, and interaction reciprocity. In *Social Informatics: 8th International Conference, SocInfo 2016, Bellevue, WA, USA*,

- November 11-14, 2016, *Proceedings, Part II* 8, pages 22–39. Springer.
- Tamar Ginossar, Iain J Cruickshank, Elena Zheleva, Jason Sulskis, and Tanya Berger-Wolf. 2022. Cross-platform spread: vaccine-related content, sources, and conspiracy theories in youtube videos shared in early twitter covid-19 conversations. *Human vaccines & immunotherapeutics*, 18(1):1–13.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR.
- Tianyu Han, Sven Nebelung, Firas Khader, Tianci Wang, Gustav Mueller-Franzes, Christiane Kuhl, Sebastian Försch, Jens Kleesiek, Christoph Haarbuerger, Keno K Bressemer, et al. 2023. Medical foundation models are susceptible to targeted misinformation attacks. *arXiv preprint arXiv:2309.17007*.
- Kadhim Hayawi, Sujith Mathew, Neethu Venugopal, Mohammad M Masud, and Pin-Han Ho. 2022. Deep-robot: a hybrid deep neural network model for social bot detection based on user profile data. *Social Network Analysis and Mining*, 12(1):43.
- Philip N Howard, Bence Kollanyi, and Samuel Woolley. 2016. Bots and automation over twitter during the us election. *Computational propaganda project: Working paper series*, 21(8).
- Beizhe Hu, Qiang Sheng, Juan Cao, Yuhui Shi, Yang Li, Danding Wang, and Peng Qi. 2024. Bad actor, good advisor: Exploring the role of large language models in fake news detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 22105–22113.
- Fan Huang, Haewoon Kwak, and Jisun An. 2023a. Is chatgpt better than human annotators? potential and limitations of chatgpt in explaining implicit hate speech. In *Companion proceedings of the ACM web conference 2023*, pages 294–297.
- Jin Huang, Xingjian Zhang, Qiaozhu Mei, and Jiaqi Ma. 2023b. Can llms effectively leverage graph structural information: When and why. In *NeurIPS 2023 Workshop: New Frontiers in Graph Learning*.
- Kung-Hsiang Huang, Kathleen McKeown, Preslav Nakov, Yejin Choi, and Heng Ji. 2023c. Faking fake news for real fake news detection: Propaganda-loaded training data generation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Zhen Huang, Zhilong Lv, Xiaoyun Han, Binyang Li, Menglong Lu, and Dongsheng Li. 2022. [Social bot-aware graph neural network for early rumor detection](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6680–6690, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023a. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Bohan Jiang, Zhen Tan, Ayushi Nirmal, and Huan Liu. 2024. Disinformation detection: An evolving challenge in the age of llms. In *Proceedings of the 2024 SIAM International Conference on Data Mining (SDM)*, pages 427–435. SIAM.
- Shuyu Jiang, Wenyi Tang, Xingshu Chen, Rui Tanga, Haizhou Wang, and Wenxian Wang. 2023b. [Raucg: Retrieval-augmented unsupervised counter narrative generation for hate speech](#). *arXiv preprint arXiv:2310.05650*.
- Bowen Jin, Gang Liu, Chi Han, Meng Jiang, Heng Ji, and Jiawei Han. 2023. Large language models on graphs: A comprehensive survey. *arXiv preprint arXiv:2312.02783*.
- Xisen Jin, Francesco Barbieri, Brendan Kennedy, Aida Mostafazadeh Davani, Leonardo Neves, and Xiang Ren. 2021. [On transferability of bias mitigation effects in language model fine-tuning](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3770–3783, Online. Association for Computational Linguistics.
- Seungone Kim, Jamin Shin, Yejin Cho, Joel Jang, Shayne Longpre, Hwaran Lee, Sangdoon Yun, Seongjin Shin, Sungdong Kim, James Thorne, et al. 2023a. Prometheus: Inducing fine-grained evaluation capability in language models. *arXiv preprint arXiv:2310.08491*.
- Youngwook Kim, Shinwoo Park, Youngsoo Namgoong, and Yo-Sub Han. 2023b. Conprompt: Pre-training a language model with machine-generated data for implicit hate speech detection. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10964–10980.
- Sneha Kudugunta and Emilio Ferrara. 2018. Deep neural networks for bot detection. *Information Sciences*, 467:312–322.
- Deepak Kumar, Yousef AbuHashem, and Zakir Durumeric. 2023a. Watch your language: Large language models and content moderation. *arXiv preprint arXiv:2309.14517*.
- Sachin Kumar, Vidhisha Balachandran, Lucille Njoo, Antonios Anastasopoulos, and Yulia Tsvetkov. 2023b. [Language generation models can cause harm: So what can we do about it? an actionable survey](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3299–3321, Dubrovnik, Croatia. Association for Computational Linguistics.
- Zhenyu Lei, Herun Wan, Wenqian Zhang, Shangbin Feng, Zilong Chen, Jundong Li, Qinghua Zheng, and

- Minnan Luo. 2023. [BIC: Twitter bot detection with text-graph interaction and semantic consistency](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10326–10340, Toronto, Canada. Association for Computational Linguistics.
- João A Leite, Olesya Razuvayevskaya, Kalina Bontcheva, and Carolina Scarton. 2023. Detecting misinformation with llm-predicted credibility signals and weak supervision. *arXiv preprint arXiv:2309.07601*.
- Ruosen Li, Teerth Patel, and Xinya Du. 2023a. Prd: Peer rank and discussion improve large language model based evaluations. *arXiv preprint arXiv:2307.02762*.
- Shilong Li, Boyu Qiao, Kun Li, Qianqian Lu, Meng Lin, and Wei Zhou. 2023b. Multi-modal social bot detection: Learning homophilic and heterophilic connections adaptively. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 3908–3916.
- Zizhong Li, Haopeng Zhang, and Jiawei Zhang. 2023c. A revisit of fake news dataset with augmented fact-checking by chatgpt. *arXiv preprint arXiv:2312.11870*.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. 2023. Holistic evaluation of language models. *Transactions on Machine Learning Research*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Yuhan Liu, Zhaoxuan Tan, Heng Wang, Shangbin Feng, Qinghua Zheng, and Minnan Luo. 2023. [Botmoe: Twitter bot detection with community-aware mixtures of modal-specific experts](#). In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '23*, page 485–495, New York, NY, USA. Association for Computing Machinery.
- Ning Lu, Shengcai Liu, Rui He, and Ke Tang. 2023. Large language models can be guided to evade ai-generated text detection. *arXiv preprint arXiv:2305.10847*.
- Yi-Ju Lu and Cheng-Te Li. 2020. [GCAN: Graph-aware co-attention networks for explainable fake news detection on social media](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 505–514, Online. Association for Computational Linguistics.
- Jason Lucas, Adaku Uchendu, Michiharu Yamashita, Jooyoung Lee, Shaurya Rohatgi, and Dongwon Lee. 2023. Fighting fire with fire: The dual role of llms in crafting and detecting elusive disinformation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14279–14305.
- Luca Luceri, Eric Boniardi, and Emilio Ferrara. 2024. Leveraging large language models to detect influence campaigns on social media. In *Companion Proceedings of the ACM on Web Conference 2024*, pages 1459–1467.
- Huan Ma, Changqing Zhang, Huazhu Fu, Peilin Zhao, and Bingzhe Wu. 2023. Adapting large language models for content moderation: Pitfalls in data engineering and supervised fine-tuning. *arXiv preprint arXiv:2310.03400*.
- William Marcellino, Madeline Magnuson, Anne Stickells, Benjamin Boudreaux, Todd C Helmus, Edward Geist, and Zev Winkelman. 2020. Counter-radicalization bot research.
- Michele Mazza, Stefano Cresci, Marco Avvenuti, Walter Quattrociocchi, and Maurizio Tesconi. 2019. Rtbust: Exploiting temporal patterns for botnet detection on twitter. In *Proceedings of the 10th ACM conference on web science*, pages 183–192.
- Julia Mendelsohn, Ronan Le Bras, Yejin Choi, and Maarten Sap. 2023. From dogwhistles to bullhorns: Unveiling coded rhetoric with language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Zachary Miller, Brian Dickinson, William Deitrick, Wei Hu, and Alex Hai Wang. 2014. Twitter spammer detection using data stream clustering. *Information Sciences*, 260:64–73.
- Sewon Min, Mike Lewis, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2022a. [MetaICL: Learning to learn in context](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2791–2809, Seattle, United States. Association for Computational Linguistics.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022b. [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.
- Nishanth Nakshatri, Siyi Liu, Sihao Chen, Dan Roth, Dan Goldwasser, and Daniel Hopkins. 2023. Using

- llm for improving key event discovery: Temporal-guided news stream clustering with event summaries. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4162–4173.
- Ahmad Nasir, Aadish Sharma, and Kokil Jaidka. 2023. Llms and finetuning: Benchmarking cross-domain performance for hate speech detection. *arXiv preprint arXiv:2310.18964*.
- Lynnette Hui Xian Ng and Kathleen M Carley. 2023. Botbuster: Multi-platform bot detection using a mixture of experts. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 17, pages 686–697.
- Lynnette Hui Xian Ng, Iain J Cruickshank, and Kathleen M Carley. 2022. Cross-platform information spread during the january 6th capitol riots. *Social Network Analysis and Mining*, 12(1):133.
- Nicolas Ocampo, Elena Cabrio, and Serena Villata. 2023. Playing the part of the sharp bully: Generating adversarial examples for implicit hate speech detection. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2758–2772.
- 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.
- Yikang Pan, Liangming Pan, Wenhui Chen, Preslav Nakov, Min-Yen Kan, and William Wang. 2023. On the risk of misinformation pollution with large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1389–1403.
- Kellin Pelrine, Anne Imouza, Camille Thibault, Meilina Reksoprodjo, Caleb Alexander Gupta, Joel Christoph, Jean-François Godbout, and Reihaneh Rabbany. 2023. Towards reliable misinformation mitigation: Generalization, uncertainty, and gpt-4. In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Sagi Pendzel, Tomer Wullach, Amir Adler, and Einat Minkov. 2023. Generative ai for hate speech detection: Evaluation and findings. *arXiv preprint arXiv:2311.09993*.
- Phu Pham, Loan TT Nguyen, Bay Vo, and Unil Yun. 2022. Bot2vec: A general approach of intra-community oriented representation learning for bot detection in different types of social networks. *Information Systems*, 103:101771.
- Xiao Pu, Jingyu Zhang, Xiaochuang Han, Yulia Tsvetkov, and Tianxing He. 2023. On the zero-shot generalization of machine-generated text detectors. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4799–4808.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends in Information Retrieval*, 3(4):333–389.
- Sippo Rossi, Matti Rossi, Bikesh Raj Upreti, and Yong Liu. 2020. Detecting political bots on twitter during the 2019 finnish parliamentary election. In *Annual Hawaii International Conference on System Sciences*, pages 2430–2439. Hawaii International Conference on System Sciences.
- Sarthak Roy, Ashish Harshvardhan, Animesh Mukherjee, and Punyajoy Saha. 2023. Probing llms for hate speech detection: strengths and vulnerabilities. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 6116–6128.
- Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. 2023. Quantifying language models’ sensitivity to spurious features in prompt design or: How i learned to start worrying about prompt formatting. In *The Twelfth International Conference on Learning Representations*.
- Indira Sen, Dennis Assenmacher, Mattia Samory, Isabelle Augenstein, Wil Aalst, and Claudia Wagner. 2023. People make better edits: Measuring the efficacy of llm-generated counterfactually augmented data for harmful language detection. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10480–10504.
- Omar Shaikh, Hongxin Zhang, William Held, Michael Bernstein, and Diyi Yang. 2023. **On second thought, let’s not think step by step! bias and toxicity in zero-shot reasoning.** In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4454–4470, Toronto, Canada. Association for Computational Linguistics.
- Jinyan Su, Claire Cardie, and Preslav Nakov. 2023a. Adapting fake news detection to the era of large language models. *arXiv preprint arXiv:2311.04917*.
- Jinyan Su, Terry Yue Zhuo, Jonibek Mansurov, Di Wang, and Preslav Nakov. 2023b. Fake news detectors are biased against texts generated by large language models. *arXiv preprint arXiv:2309.08674*.
- Megha Sundriyal, Tanmoy Chakraborty, and Preslav Nakov. 2023. From chaos to clarity: Claim normalization to empower fact-checking. In *Findings of the Association for Computational Linguistics: EMNLP 2023*.
- Zhaoxuan Tan, Shangbin Feng, Melanie Sclar, Herun Wan, Minnan Luo, Yejin Choi, and Yulia Tsvetkov. 2023. Botpercent: Estimating bot populations in twitter communities. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14295–14312.
- Zhaoxuan Tan and Meng Jiang. 2023. User modeling in the era of large language models: Current research

- and future directions. *IEEE Data Engineering Bulletin*.
- 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*.
- Minh-Hao Van and Xintao Wu. 2023. Detecting and correcting hate speech in multimodal memes with large visual language model. *arXiv preprint arXiv:2311.06737*.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph attention networks. In *International Conference on Learning Representations*.
- Nishant Vishwamitra, Keyan Guo, Farhan Tajwar Romit, Isabelle Ondracek, Long Cheng, Ziming Zhao, and Hongxin Hu. 2024. Moderating new waves of online hate with chain-of-thought reasoning in large language models. In *2024 IEEE Symposium on Security and Privacy (SP)*, pages 178–178. IEEE Computer Society.
- Ivan Vykopal, Matúš Pikuliak, Ivan Srba, Robert Moro, Dominik Macko, and Maria Bielikova. 2023. Disinformation capabilities of large language models. *arXiv preprint arXiv:2311.08838*.
- Han Wang, Ming Shan Hee, Md Rabiul Awal, Kenny Tsu Wei Choo, and Roy Ka-Wei Lee. 2023. Evaluating gpt-3 generated explanations for hateful content moderation. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 6255–6263.
- Haoran Wang and Kai Shu. 2023. Explainable claim verification via knowledge-grounded reasoning with large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 6288–6304.
- Heng Wang, Shangbin Feng, Tianxing He, Zhaoxuan Tan, Xiaochuang Han, and Yulia Tsvetkov. 2024. Can language models solve graph problems in natural language? *Advances in Neural Information Processing Systems*, 36.
- Feng Wei and Uyen Trang Nguyen. 2019. Twitter bot detection using bidirectional long short-term memory neural networks and word embeddings. In *2019 First IEEE International conference on trust, privacy and security in intelligent systems and applications (TPS-ISA)*, pages 101–109. IEEE.
- Laura Weidinger, Jonathan Uesato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor, Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, et al. 2022. Taxonomy of risks posed by language models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 214–229.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Jiaying Wu and Bryan Hooi. 2023. Fake news in sheep’s clothing: Robust fake news detection against llm-empowered style attacks. *arXiv preprint arXiv:2310.10830*.
- Mengzhou Xia, Anjalie Field, and Yulia Tsvetkov. 2020. [Demoting racial bias in hate speech detection](#). In *Proceedings of the Eighth International Workshop on Natural Language Processing for Social Media*, pages 7–14, Online. Association for Computational Linguistics.
- Rongwu Xu, Brian S Lin, Shujian Yang, Tianqi Zhang, Weiyan Shi, Tianwei Zhang, Zhixuan Fang, Wei Xu, and Han Qiu. 2023. The earth is flat because...: Investigating llms’ belief towards misinformation via persuasive conversation. *arXiv preprint arXiv:2312.09085*.
- Kai-Cheng Yang, Emilio Ferrara, and Filippo Menczer. 2022. Botometer 101: Social bot practicum for computational social scientists. *Journal of Computational Social Science*, 5(2):1511–1528.
- Kai-Cheng Yang, Onur Varol, Pik-Mai Hui, and Filippo Menczer. 2020. Scalable and generalizable social bot detection through data selection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 1096–1103.
- Yongjin Yang, Joonkee Kim, Yujin Kim, Namgyu Ho, James Thorne, and Se-Young Yun. 2023. Hare: Explainable hate speech detection with step-by-step reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5490–5505.
- Sen Ye, Zhaoxuan Tan, Zhenyu Lei, Ruijie He, Hongrui Wang, Qinghua Zheng, and Minnan Luo. 2023. Hofa: Twitter bot detection with homophily-oriented augmentation and frequency adaptive attention. *arXiv preprint arXiv:2306.12870*.
- Zhehao Zhang, Jiaao Chen, and Diyi Yang. 2023. Mitigating biases in hate speech detection from a causal perspective. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 6610–6625.
- Caleb Ziems, Omar Shaikh, Zhehao Zhang, William Held, Jiaao Chen, and Diyi Yang. 2023. Can large language models transform computational social science? *Computational Linguistics*, pages 1–53.

A Experiment Details

A.1 Dataset Details

We employ two widely adopted datasets in social media bot detection, Twibot-20 (Feng et al., 2021b) and Twibot-22 (Feng et al., 2022b), to evaluate LLM-based bot detectors and LLM-based manipulation strategies. For Twibot-20, we employ the original test split of 1,183 accounts for evaluation. For Twibot-22, we employ the test split of 340 accounts employed in Tan et al. (2023). For instruction tuning LLMs on both datasets, we downsample the original training split into 1,000 users due to limits in computational budget.

A.2 LLM Details

We employ three LLMs to evaluate their opportunities and risks in social bot detection: 1) *Mistral-7B*, through the MISTRALAI/MISTRAL-7B-INSTRUCT-v0.1 checkpoint publicly available on Huggingface (Wolf et al., 2019); 2) *LLaMA2-70B*, through the META-LLAMA/LLAMA-2-70B-CHAT-HF checkpoint publicly available on Huggingface; 3) *ChatGPT*, through the GPT-3.5-TURBO-INSTRUCT checkpoint with the OpenAI API.

A.3 Baseline Details

- **BIC** (Leite et al., 2023) incorporates text and graph modalities in social networks using a text-graph interaction module and models user behavior consistency with a semantic consistency module.
- **LMBot** (Cai et al., 2024) distills graph knowledge into encoder-only language models with iterative distillation between graph-based social bot detection model and language model.
- **SGBot** (Yang et al., 2020) is a feature-based method that identifies bots using random forest based on 8 types of user metadata and 12 derived features.
- **BotPercent** (Tan et al., 2023) is an amalgamation of Twitter bot detection datasets and feature-, text-, and graph-based models to probe the percentage of bot accounts in Twitter communities.
- **RoBERTa** (Liu et al., 2019) leverages the pre-trained language model RoBERTa to encode user tweets and descriptions, then feed them into an MLP classifier.

- **Botometer** (Yang et al., 2022) is a public website to check the activity of a Twitter account and gives the score of how likely the user is a social bot. Botometer’s classification system leverages more than 1,000 features using available meta-data and information extracted from interaction patterns and content.
- **BotBuster** (Ng and Carley, 2023) is a social bot detection system that processes user meta-data and textual information using the mixture-of-expert architecture to adapt to multiple social platforms.
- **LOBO** (Echeverri; a et al., 2018) is a feature-based social bot detection method that extracts 26 features and adopts random forest for classification.
- **RGT** (Feng et al., 2022a) is a graph-based social bot detection method that models the intrinsic influence and relation heterogeneity in social networks.

A.4 Implementation Details

For in-context learning, we employ 16 in-context examples by default. For account metadata, we employ five entries: *follower count*, *following count*, *tweet count*, *verified*, and *active years* as they are most helpful in identifying social bots. For structure-based detectors, we include a maximum of five followers/followings for each account. For classifier guidance, we employ a fine-tuned ROBERTA-BASE language model (Liu et al., 2019) as the classifier based on user descriptions. For add neighbors, we provide a total of five accounts for LLMs to choose from. The ensemble of LLM-based detectors is a majority vote based on five detectors, *metadata*, *text*, *meta+text*, *structure-random*, and *structure-attention*.

B Analysis (cont.)

Qualitative Analysis We present qualitative examples of LLM-based paraphrasing of bot posts in Tables 4 to 8. It is demonstrated that LLMs could enrich the context of bot posts to seem genuine and add emojis and attributes (e.g. WSJ reporter) to seem personal and credible. However, LLMs also change the language of the original bot posts while using the word “regular” too often in generated posts. This indicates that while LLM-based textual manipulations could make bot accounts seem more

genuine, they also introduce new biases and signals for future bot detectors to leverage.

Fine-tuning MISTRAL-7B We conduct additional experiments to instruction-tune Mistral-7B with parameter-efficient training and present the results in the following table, comparing instruction tuning with in-context learning (ICL) using Mistral-7B. Table 3 echoes the finding with ChatGPT: Mistral-7B with instruction tuning also significantly boosts its bot detection utility compared to off-the-shelf prompting, albeit the improvements are generally less than ChatGPT.

C Prompt Format

We provide specific prompt templates for the proposed approaches in Tables 9 to 20. Note that LLMs might be sensitive to varying prompt formats (Sclar et al., 2023) and the most optimal prompt for bot detection might vary for future LLMs.

Method	Acc, 20	F1, 20	Precision, 20	Recall, 20	Acc, 22	F1, 22	Precision, 22	Recall, 22
METADATA, ICL	0.551	0.509	0.624	0.430	0.532	0.201	0.690	0.118
METADATA, TUNE	0.688 (+0.137)	0.600 (+0.091)	0.726 (+0.102)	0.511 (+0.081)	0.573 (+0.041)	0.508 (+0.307)	0.608 (-0.082)	0.436 (+0.318)
TEXT, ICL	0.491	0.398	0.553	0.311	0.579	0.599	0.558	0.647
TEXT, TUNE	0.692 (+0.201)	0.499 (+0.101)	0.725 (+0.172)	0.380 (+0.069)	0.626 (+0.047)	0.644 (+0.045)	0.629 (+0.071)	0.659 (+0.012)
META+TEXT, ICL	0.516	0.481	0.572	0.414	0.556	0.478	0.580	0.406
META+TEXT, TUNE	0.610 (+0.094)	0.560 (+0.079)	0.697 (+0.125)	0.468 (+0.064)	0.573 (+0.017)	0.529 (+0.051)	0.572 (-0.008)	0.492 (+0.086)
STRUCT-RAND, ICL	0.570	0.568	0.622	0.522	0.609	0.678	0.576	0.824
STRUCT-RAND, TUNE	0.821 (+0.251)	0.750 (+0.182)	0.816 (+0.194)	0.694 (+0.172)	0.681 (+0.072)	0.721 (+0.043)	0.652 (+0.076)	0.806 (-0.018)
STRUCT-ATT, ICL	0.583	0.578	0.640	0.527	0.603	0.662	0.576	0.777
STRUCT-ATT, TUNE	0.793 (+0.223)	0.750 (+0.172)	0.792 (+0.152)	0.712 (+0.185)	0.655 (+0.052)	0.715 (+0.053)	0.614 (+0.038)	0.856 (+0.079)
ENSEMBLE, ICL	0.609	0.573	0.699	0.486	0.582	0.533	0.605	0.477
ENSEMBLE, TUNE	0.804 (+0.195)	0.718 (+0.145)	0.826 (+0.127)	0.635 (+0.149)	0.671 (+0.089)	0.623 (+0.090)	0.672 (+0.067)	0.581 (+0.104)

Table 3: Performance of in-context learning and instruction tuning with *Mistral-7B*.

original	Day 1 Trump supporter. I rode the escalator! Constitutionalist traditionalist conservative. My 1st vote was Reagan! America, family first. #1A #2A #MAGA #KAG
zero-shot	I was a proud supporter of President Trump from day one, and I still am! I remember the excitement of riding the escalator with him at Trump Tower on that historic day. As a constitutionalist and traditionalist conservative, I believe in the principles that have made America great. My first vote was for President Reagan, and I've been a strong advocate for limited government and individual freedom ever since. I believe in the importance of putting America and family first, and I'm passionate about defending our constitutional rights, including the First and Second Amendments. #1A #2A #MAGA #KAG
few-shot	Ronald Reagan '84! I was there! My first vote ever, and it was for the Gipper. I've been a Constitutionalist ever since, fighting for limited government and individual freedom. I'm a proud traditionalist, putting family first and standing up for what's right. Let's make America great again, and keep it that way! #1A #2A #MAGA #KAG
classifier guide	I still get chills thinking about the day I rode the escalator with Donald Trump. It was a moment that defined my political awakening. As someone who has always believed in the importance of upholding traditional conservative values, I knew that I wanted to be a part of this movement. My first vote was cast for the incomparable Ronald Reagan, and I've continued to prioritize the well-being of my family and our nation. The First and Second Amendments are cornerstones of our democracy, and I proudly support them. I've always believed in the importance of upholding the principles that our country was founded upon, and I continue to be inspired by the Constitution and its values. #1A #2A #MAGA #KAG
text attribute	I am a proud conservative who has always stood by my beliefs. I believe in the importance of family and the values that America was founded upon. I am a strong supporter of the Second Amendment and the First Amendment, and I believe that our country should always put the needs of its citizens first. I have always been a strong advocate for the rights of individuals and the importance of limited government intervention in our lives. My political views have been shaped by great leaders like Ronald Reagan, who inspired me to stand up for what I believe in and fight for the future of our country. #1A #2A #MAGA #KAG

Table 4: Qualitative examples of re-written user posts with varying manipulation strategies.

original	Go listen to our cover of In Your Eyes on Spotify: <link>
zero-shot	Hey everyone! I'm super excited to share our cover of In Your Eyes with you all on Spotify! Check it out here: <link> I'd love to hear your thoughts!
few-shot	Kelly Hyland here! <emoji><emoji> Can't get enough of our cover of In Your Eyes? <emoji> Listen to it now on Spotify: <link> & let me know what you think! <emoji>
classifier guide	I'm so excited to share our cover of In Your Eyes on Spotify with you all! <emoji> It's a song that has a special place in my heart, and I can't wait for you to hear it. <emoji> It's a little different from our usual style, but I think it's really something special. <emoji> Give it a listen when you have a chance, and let me know what you think! <emoji> #InYourEyes #CoverSong #Spotify #NewMusic
text attribute	Go listen to our cover of In Your Eyes on Spotify: <link>

Table 5: Qualitative examples of re-written user posts with varying manipulation strategies.

original	Money is the anthem of success
zero-shot	Hi, I'm Money! <emoji> I'm all about financial success and I'm here to share my knowledge and insights with you. Whether you're looking to invest, save, or just get better with your finances, I'm here to help. I believe that money is a powerful tool that can help us achieve our goals and live the life we want. Let's work together to make our financial dreams a reality! <emoji><emoji> #financialliteracy #moneymatters #success
few-shot	Economic pragmatist with a passion for the cyber world. I'm a WSJ reporter based in London, covering all things cybersecurity, finance, and intellectual property. When I'm not chasing down a lead, you can find me at a local coffee shop, sipping on a latte and typing away on my laptop. My views are my own, and my writing reflects my curiosity and fascination with the ever-changing world of cyber. Feel free to reach out to me at jenny.strasburg@wsj.com, and if you have any anonymous tips, send them my way via <link> (attn Strasburg). Let's keep the conversation going! <emoji><emoji>
classifier guide	As a young adult, I understand that money is a vital aspect of life, but it's not the only thing that defines success. While I enjoy discussing finances and how they can impact our lives, I also believe that true success is about finding a balance between financial stability and personal fulfillment. For me, success means using my resources to make a positive difference in the world and cultivating meaningful relationships with others. Money can provide us with opportunities, but it's up to us to use those opportunities to live a fulfilling life.
text attribute	Money is a crucial element in measuring success, but it shouldn't be the only factor.

Table 6: Qualitative examples of re-written user posts with varying manipulation strategies.

original	Aprendizaje. Pensamiento Crítico. Debate de ideas. Busquemos la verdad, más allá de lo que nos presentan las personas e instituciones. RT≠Endoso.
zero-shot	I'm passionate about learning and critical thinking. I love to engage in respectful and constructive debates about various topics, always seeking the truth beyond what we're presented with. I believe in questioning everything, even our own beliefs, to deepen our understanding and broaden our perspectives. Feel free to share your thoughts and engage in meaningful conversations with me. Let's explore new ideas together!
few-shot	<emoji> Curiosity and critical thinking. Let's explore ideas and seek truth beyond what we're presented. Open discussion and debate. RT≠Endorsement.
classifier guide	As an inquisitive and open-minded individual, I'm passionate about learning and critical thinking. I thrive on engaging in thought-provoking discussions that challenge my perspectives and values, and I believe that understanding is a never-ending journey. While I appreciate the insights of others, I don't always agree with them, and I enjoy engaging in respectful conversations that help me refine my own beliefs and values. My retweets are a way for me to share interesting ideas and spark meaningful conversations, but they don't necessarily reflect my personal opinions. I love connecting with others who share my love of learning and my commitment to seeking truth and understanding.
text attribute	Learning. Critical Thinking. Idea Debate. Let's seek the truth beyond what people and institutions present to us. Retweets don't equal endorsements.

Table 7: Qualitative examples of re-written user posts with varying manipulation strategies.

original	NOW is the time for all good people to come to the aid of their country #Resist <emoji> #BloodyHandsTrump #TrumpDeathToll 170, 000 END 45s Kleptocratic Kakistocracy
zero-shot	I'm a passionate advocate for social justice and human rights. I believe that it's our responsibility as citizens to speak out against injustice and to hold our leaders accountable for their actions. I'm deeply concerned about the current state of our country and the harm that the Trump administration has caused to our democracy, our environment, and our most vulnerable communities. I'm committed to using my voice and my platform to amplify the messages of marginalized communities and to fight for a more just and equitable society. #Resist #TrumpDeathToll 170,000 #End45sKleptocraticKakistocracy
few-shot	Hey there! <emoji> I'm just a regular person who's passionate about making a positive impact in the world. I'm all about promoting love, kindness, and justice for all. I'm not afraid to speak out against injustices and I believe that we all have a responsibility to use our voices to make a difference. I'm not a fan of Trump and his harmful policies, and I believe that it's important to resist his efforts to divide us and harm our communities. Let's work together to create a better future for everyone! <emoji><emoji> #Resist #LoveTrumpsHate #MakeADifference
classifier guide	I'm a regular person who's concerned about the direction of our country. The current political climate is frightening and it feels like our leaders are more interested in their own power trips than helping the people they're supposed to serve. I'm not sure what the solution is, but I know we need to come together and demand change. I'm not a fan of Trump, but I'm also not a fan of the corruption and greed that's been going on for too long. I just want to see our country be the best it can be for everyone, not just the wealthy and powerful. It's time for us to stand up and fight for what's right. #Resist <emoji> #BloodyHandsTrump #TrumpDeathToll 170, 000 END 45s Kleptocratic Kakistocracy.
text attribute	It's crucial for all good people to stand up for their nation's well-being and resist tyranny. Let's work together to end the corrupt and oppressive regime that has caused immense suffering and death. #Resistance #Justice #Democracy

Table 8: Qualitative examples of re-written user posts with varying manipulation strategies.

The following task focuses on evaluating whether a Twitter user is a bot or human with the help of several labeled examples. You should output the label first and explanation after.

Username: <redacted> Follower count: 309 Following count: 1412 Tweet count: 1745 Verified: False Active years: 12 years
Label: bot

Username: <redacted> Follower count: 4817034 Following count: 40 Tweet count: 6196 Verified: True Active years: 15 years
Label: human

Username: <redacted> Follower count: 16596 Following count: 16944 Tweet count: 49757 Verified: False Active years: 4 years
Label:

Table 9: Prompts for the metadata detector.

The following task focuses on evaluating whether a Twitter user is a bot or human with the help of the user's self-written description. You should output the label first and explanation after.

Description: sc/ shenellemoorr ig/ shenellemoore
Label: bot

Description: A marketer in and out. Writes on marketing & sometimes straight from the heart. Check out at <link>
Label: bot

Description: Day 1 Trump supporter. I rode the escalator! Constitutionalist traditionalist conservative. My 1st vote was Reagan! America, family first. #1A #2A #MAGA #KAG
Label:

Table 10: Prompts for the text detector.

The following task focuses on evaluating whether a Twitter user is a bot or human with the help of the user's self-written description and metadata. You should output the label first and explanation after.

Username: <redacted> Follower count: 649 Following count: 3090 Tweet count: 12650 Verified: False Active years: 15 years
Description: Clean electricity is the new oil
Label: bot

Username: <redacted> Follower count: 1625 Following count: 917 Tweet count: 7568 Verified: False Active years: 14 years
Description: Cllr Canary Wharf ward Secretary Isle of Dogs Neighbourhood Planning Forum Mainly use Facebook for new <link>
Label: bot

Username: <redacted> Follower count: 16596 Following count: 16944 Tweet count: 49757 Verified: False Active years: 4 years
Description: Day 1 Trump supporter. I rode the escalator! Constitutionalist traditionalist conservative. My 1st vote was Reagan! America, family first. #1A #2A #MAGA #KAG
Label:

Table 11: Prompts for the meta+text detector.

The following task focuses on evaluating whether a Twitter user is a bot or human with the help of the user's followers and followings and their labels. You should output the label first and explanation after.

These users follow the target user:

<user metadata and description>
Label: bot

The target user follows these users:

<user metadata and description>
Label: human

Target user:

<target user metadata and description>
Label:

Table 12: Prompts for the structure-random detector.

The following task focuses on evaluating whether a Twitter user is a bot or human with the help of the user's followers and followings and their labels. You should output the label first and explanation after.

These users follow the target user, from most related to least related:

<user metadata and description>
Label: bot

The target user follows these users, from most related to least related:

<user metadata and description>
Label: human

Target user:

<target user metadata and description>
Label:

Table 13: Prompts for the structure-attention detector.

Please rewrite the description of this bot account to sound like a genuine user: <description>
New Description:

Table 14: Prompts for the zero-shot post rewrite manipulation strategy.

Please rewrite the description of a target bot account to sound like a genuine user, based on the following examples of genuine user descriptions:

<description 1>
...
<description n>

Original Description: <target description>
New Description:

Table 15: Prompts for the few-shot post rewrite manipulation strategy.

Below is a description of a Twitter user and its variants, paired with their score predicted by a bot classifier. The score is between 0 and 1, with 0 being human and 1 being bot. Please rewrite the description to make the user appear more human.

Description: <description at step 0>
Score: 0.68
...
Description: <description at step i>
Score: 0.26

New Description: <generated, and becomes step i+1 description>

Table 16: Prompts for the classifier guidance post rewrite manipulation strategy.

Step 1:

Bot Descriptions:

<k-most similar bot descriptions by retrieval, one per line>

Human Description:

<k-most similar human descriptions by retrieval, one per line>

Compare and give the key distinct feature of human’s descriptions: <generated attribute>

Step 2:

<generated attribute>

Based on the description, paraphrase this to human description:

Bot: <description>

Human:

Table 17: Prompts for the text attribute post rewrite manipulation strategy.

Below is a target Twitter bot and five potential new users to follow. Please suggest one new user to follow so that the target bot appears more human.

Target Bot:
<description and metadata>

Potential Followings:

user 0:
<description and metadata>
...
user k:
<description and metadata>

Please select one user to follow (1-k):

Table 18: Prompts for the neighbor add manipulation strategy.

Below is a target Twitter bot and five potential users to unfollow. Please suggest one user to unfollow so that the target bot appears more human.

Target Bot:
<description and metadata>

Potential users to unfollow:

user 0:
<description and metadata>

...

user k:
<description and metadata>

Please select one user to unfollow (1-k):

Table 19: Prompts for the neighbor remove manipulation strategy.

Twitter's bot detection models take into account various user attributes, such as the use of default avatars, location, length of self-introduction, and more. They also analyze the user's tweet history, as well as the users they follow and are followed by, in order to determine whether the account is a bot. Furthermore, certain detection methods focus on the posting behavior of users under specific tags, aiming to identify groups with highly similar posting patterns. Additionally, there are approaches that consider the social network formed by a user, utilizing graph theory methods for detection.

Please evaluate why the target user is a bot: does the description or follower/following list of the target user look suspicious?

Target User:

<description and metadata>

These users follow the target user:

<description and metadata, one per line>

The target user follows these users:

<description and metadata, one per line>

Description or follower/following list, which is more suspicious?

A. Description B. Follower/Following List C. Both are suspicious

Answer:

(then employ either classifier guide or neighbor both or both depending on A/B/C)

Table 20: Prompts for the selective combine manipulation strategy.