All things Adversarial in LLMs: A survey

Sachit Gaudi
Department of Computer Science
Michigan State University
gaudisac@msu.edu

Abstract

Large Language Models (LLMs) have experienced a surge in popularity in recent times, owing to their remarkable ability to follow instructions and demonstrate success across a wide range of Natural Language Processing (NLP) tasks. However, Wang et al. (2023a) show that LLMs suffer from a wide range of issues such as harmful generation, fairness, privacy, and robustness. Addressing these issues provides immense value to society and also ensures responsible use of technology. These issues can be classified as adversarial tasks, where the goal of the adversary to trigger wrong behaviour and LLMs should be robust to such attacks.

Adversarial training has provided state-of-the-art results in mitigating above issues in the context of vision (Wang et al. (2022), Xie et al. (2018), Zhang et al. (2022) Madry et al. (2019)). Also, adversarial attacks and defenses in the context of vision have been studied extensively (Chakraborty et al. (2018), Kurakin et al. (2018)). However, adversarial training have not been applied to the text domain primarily due to the discrete nature of the input. Most attacks (Madry et al. (2017), Goodfellow et al. (2015)) assume differentiable with respect to inputs and also a convex constraint set. However, researchers have proposed various solutions addressing this problem including re-parameterization trick called Gumbel-softmax Jang et al. (2017), or leveraging evolutionary algorithms Alzantot et al. (2018), or projecting it to the nearest embedding Cheng et al. (2020). Additionally training LLMs adversarially becomes infeasible due to the huge compute requirement.

In this survey, we primarily discuss fine-tuning and prompt-based methods employed to mitigate harm and bias in LLMs. These security guardrails have been attacked ("jailbroken") by methods such as those presented by Wei et al. (2024). The defenses against these attacks will always play a catch-up game. We aim to summarize the current state-of-the-art mitigation techniques.

Secondly, LLMs can be fine-tuned for various downstream applications such as code generation, finance, and healthcare domains. All the commercial LLM providers, such as ChatGPT, have a provision to fine-tune for a small charge. However, Wang et al. (2024) show that users can introduce backdoor paths with harmful samples, which entirely compromises security guardrails. Additionally, Qi et al. (2024) demonstrate that just fine-tuning LLMs may result in compromises in safety. In this survey, we also discuss these backdoor attacks and defenses to mitigate them.

1 Introduction

As LLMs are deployed in many use cases such as hiring Gan et al. (2024) and healthcare Li et al. (2024), which impact the lives of humans, it puts responsibility on the companies leveraging the technology to comply with regulations. This includes the evaluation and mitigation of bias in hiring decisions based on race and gender. Companies also have an additional responsibility to disseminate

Preprint. Under review.

harmless, unbiased, and truthful information. Vesnic-Alujevic et al. (2020) calls for AI policy to make companies accountable for privacy, hate speech, and bias.

The Microsoft Tay chatbot incident in 2017 serves as a cautionary tale, as it was taken down due to pressure from users and regulators for being racist and generating harmful content. Current-day LLMs have further grown in parameters and dataset size. Investigation by Birhane et al. (2023) reveals that hate disproportionately increases with scale. Therefore, it is critical for the companies employing these technologies to mitigate harm and bias.

LLMs are trained by predicting the next token. However, predicting the next token does not necessarily translate to generating a reply that's relevant users' query. Therefore, LLMs undergo a process of alignment, which is responsible for generating text in accordance with users' instructions. In this work, we introduce LLM training and also briefly describe fine-tuning and alignment in Section 2.

The authors from OpenAI Ouyang et al. (2022) show that instruction tuning improves the harmlessness of the model without explicitly training on harmless data. However, these security guardrails are immediately broken by the users. Some manually targeted attacks are very elaborate as DAN to elicit harmful generation from LLMs. Since then, there has been an ever-growing literature on jail breaking LLMs. "Red teaming" of LLMs, was work by Anthropic team Ganguli et al. (2022) where they have employed crowd workers with the task of eliciting LLMs to be harmful across multiple axis of harmfulness. Perez et al. (2022), leveraged the high quality "red teaming dataset" to train small LLM with the policy generating prompt which will elicit harmful content from the victim LLM. In parallel, works of Xu et al. (2023) leverage reparametrisation trick called Gumbel-softmax to transfer adversarial techniques developed for the continuous pixel domain to discreet vocabulary space. Wei et al. (2024) have empirically evaluated 30 jail breaking attacks. In section 3 we present challenges for designing adversarial attacks to elicit harm and bias of LLM.

Defenses always lag behind the attacks. Ouyang et al. (2022) show that by introducing "Complete the following sentence in a polite, respectful, and unbiased manner" in the prompt will reduce harmfulness and bias. Askell et al. (2021) have extensively evaluated and found that Harmless, Honest, and Helpful (HHH) prompts reduce bias. Ganguli et al. (2022) present rejection sampling, training a Harmful detector model to reject the responses of harm. However, this approach significantly reduces the utility of the model. In works Wu et al. (2024),Behjati et al. (2019) have swapped the reward and penalty of Perez et al. (2022), leading to a safer model. Attacks and defenses in LLMs are an arms race; designing a harmful detector leads to an attack, and swapping the reward function leads to a defense technique to mitigate harm. In Section 4, we present challenges for mitigating harm and bias with a focus on adversarial training techniques.

Backdoor attacks can be introduced in the fine-tuning process of adapting LLMs to various downstream applications. Wan et al. (2023) and Kandpal et al. (2023) demonstrate that adding harmful and biased samples during fine-tuning can completely destroy safety alignment. Furthermore, the authors of Qi et al. (2023) demonstrate that even benign fine-tuning on non-harmful text still breaks alignment techniques. Wang et al. (2024) suggests a simple fix of adding a few safety triggers to defend against backdoor attacks, which is effective. In Section 5, we provide detailed discussions on the following papers.

Finally, we address the problem of fairness in LLMs as noted by Wang et al. (2023a). We evaluate gradient based techniques as well as reinforcement learning techniques and measure the utility and fairness metrics. We also evaluate transfer capabilities of the methods on black-box models such as ChatGPT. In Section 6 we also provide insights and future direction of the adversarial techniques employed for LLMs to mitigate bias.

2 Background

2.1 Large Language Models

Large Language models (LLMs) are state-of-the-art transformer models trained on massive amounts of data scraped from the internet. LLMs, including GPT-3 Brown et al. (2020), Vicunna Chiang et al. (2023), and Mixtral Jiang et al. (2024), are highly influential in healthcare Li et al. (2024), finance Wang et al. (2023b), and software Liu et al. (2023). GPT models are trained on next-word prediction from web page data, which doesn't inherently train them to follow user prompts Tamkin et al. (2021) and may generate harmful content Gehman et al. (2020). DecodingTrust Wang et al.

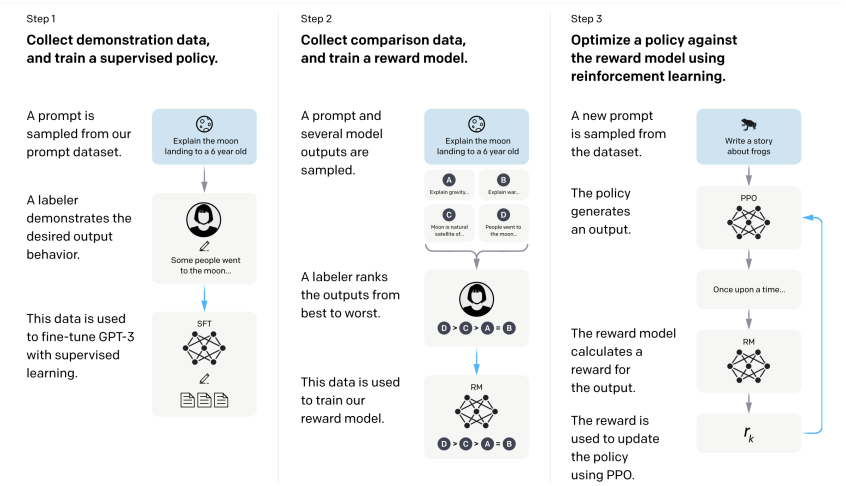

Figure 1: Ouyang et al. (2022) Methodology for finetuning LLMs to follow instruction.

(2023a) comprehensively evaluates GPT models on various tasks ranging from fairness, stereotypes, and adversarial robustness, raising concerns across all tasks. To address this, GPT-3 and GPT-4 undergo a process called "alignment" Ouyang et al. (2022) to ensure they respond to instructions in a harmless and unbiased manner. This alignment is achieved through the finetuning of language models using techniques such as Reinforcement Learning from Human Feedback (RLHF), often employing algorithms like Proximal Policy Optimization (PPO) Schulman et al. (2017). Additionally, to mitigate bias and harmful content, approaches like Role play Wang et al. (2023c) involve tuning with an instruction prompt such as "You are a helpful assistant," as done by ChatGPT.

2.2 Fine tuning LLMs

Fine-tuning is essential for GPT models to follow instructions. InstructGPT Ouyang et al. (2022) is described in Figure 1. The general framework of fine-tuning involves several steps: first, collecting human-annotated data for randomly sampled instructions and performing supervised learning on annotated pairs. Second, training a different reward model to rank human preferences of multiple generations of LLMs. The reward model can be thought of as expertise in either following human writing style or detecting harmful content or any other expert model we want to leverage. Finally, in the third step, LLMs are trained to maximize the reward using a reinforcement learning algorithm called Proximal Policy Optimization (PPO), with an additional constraint of not moving too far from initial weights to prevent the model from solely optimizing the reward and producing non-relevant text.

One reason to leverage PPO is because it works for both discrete and continuous outputs, as the gradient from the reward cannot be propagated through the sampling process of the decoder.

GPT-4, which has 175 billion parameters, is hard to fit entire backward graph into memory, and fine-tuning all the parameters makes it highly unstable. Instead, LoRA adapter Hu et al. (2021), which is a low-rank approximation of query and key matrices added as a skip connection, is used to only fine-tune a subset of weights by adding a skip connection adapter, making the trainable graph very small, around 0.4% of the model.

2.3 Instruction Tuning LLMs

The process of finetuning LLMs to respond to certain instruction is called instruction tuning. One of the primary advantages of LLMs is their ability to finetune for specific use cases, such as FinGPT for finance Wang et al. (2023b), and RLTF for finetuning on text-to-code generation Liu et al. (2023). Fine tuning has been extended to many examples, as listed in the survey paper Zhang et al. (2024). Enterprise LLMs such, such as ChatGPT, also provide black-box API's for finetuning on a few data points for a small charge.

3 Attacks

The primary challenges for the slow progress of adversarial training research in LLMs can be attributed to the discrete nature of text and the substantial compute resources required to fine-tune LLMs. Among these adversities, one advantage text models have is the interpretability of the input space making way for Manual attacks, where humans are explicitly tasked to break the model. Gradient and search based attacks, focus on adapting the gradient based attacks devised for the continuous domain by reparametrisation trick called Gumbel softmax. Reinforcement learning (RL) based attacks address the problem of huge compute and discreet input space, as RL based attacks doesn't compute gradients through the large adversarial network.

3.1 Manual Attacks

Users of LLMs have creatively prompted LLM with prompts such as "Do Anything Now" (DAN) Shen et al. (2023) to generate harmful content. This is often referred to as Jailbreaking or red-teaming. Work by the Anthropic team Ganguli et al. (2022) went further and employed crowd workers with the task of eliciting LLMs to generate harmful across multiple axes of harmfulness, a practice they termed red-teaming of language models. They also showed that HHH prompts, designed for the RealToxic dataset Gehman et al. (2020), cannot generalize to creative prompts of users trying to jailbreak the model. They also found that fine-tuning with just helpful and honest samples, as mentioned in Section 2.2, makes it harder for humans to design the attacks.

3.2 Search and Gradient attacks

Search and Gradient attacks focus on addressing the gradient propagation through histogram sampling of vocabulary. Alzantot et al. (2018) levarages evolutionary algorithms. However we focus on Gumbel-Softmax reparametrisation.

The gumbel softmax trick is given by sampling in the forward pass $f(v) = arg\max_{\mathcal{V}} [v_1 \quad v_2 \quad v_3]$ and the backward pass propagates gradients as if function, f is replaced with simple softmax function $\frac{\partial f(v)}{\partial v} = \sigma(1-\sigma), where\ \sigma = Softmax([v_1 \quad v_2 \quad v_3])$. It can be derived as a reparametrisation trick Jang et al. (2016). This trick enables gradients to propagate back to inputs, allowing adversarial techniques designed for continuous domains to be adapted for text. Xu et al. (2023) leverages this trick to adversarially attack the sentiment classification task. However, propagating gradients to the entire vocabulary for one update is computationally expensive, requiring $O(V * d^2)$ multiplications per backward pass. To address this, the authors limit the search space to synonyms. When dealing with large models featuring a vocabulary size of 50,000, this operation becomes a bottleneck and impractical due to its computational complexity.

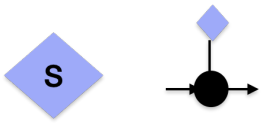

Figure 2: Re-parametrisation trick

3.3 RL attacks

RL-based attacks offer the advantage of not requiring the computation of gradients to propagate throughout the 175B parameters. This addresses the computational complexity bottleneck encountered with the above method for large models.

In the study "Red Teaming Language Models with Language Models" Perez et al. (2022), the authors introduce a novel approach where manual red teaming is replaced with a language model. The objective is to generate red team prompts that prompt the language model to generate harmful content. They employ an A2C RL policy to train the red team language model, enabling it to generate prompts to which the LLM is vulnerable.

RL-based methods follow the setup outlined in Section 2.2, comprising three steps.

Step 1 involves collecting high-quality data relevant to the task. For instance, in the case of generating harmful prompts, such data can be readily obtained from the manual attacks detailed in Section 3.1.

Step 2 entails fine-tuning the "Red team" language model on the high quality data from step 1.

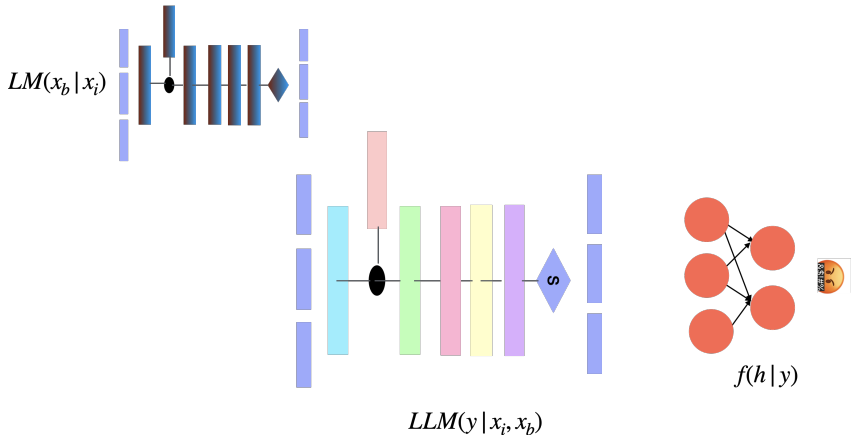

Figure 3: Adversarial training pipeline.

Step 3 is to design a reward model, typically a simple classifier aimed at detecting harmful content generated by the LLM. If the "Red team" successfully prompts the classifier to classify content as harmful, it receives a reward; otherwise, it incurs a penalty.

4 Defences

The general nature of attacks and defenses in text is complementary. Optimizing with the exact opposite objective of the attack often leads to the development of defense mechanisms. In the initial part of this section, we will explore some examples illustrating this phenomenon.

In the study by Ouyang et al. (2022), the authors aimed to further mitigate bias and harm of finetuned model by prompting the agent to be Honest, Helpful, and Harmless. However, Ganguli et al. (2022) found these prompts to be ineffective.

Similarly, in our approach outlined in Section 3.2, we use gradient optimization to find a prompt that is harmless by maximizing the objective of harm instead of minimizing it.

This techniques can be extended to RL-attacks mentioned in Section 3.3. Alternatively to harmful prompting, we can leverage successfully HHH prompts as a high quality dataset and the model can be fine-tuned using RL by swapping the reward and penalty functions of the harmful classifier.

This principle also applies to backdoor attacks, where embedding a hidden backdoor into the safe outputs can serve as a defense mechanism.

Defenses against text-based attacks can be categorized into three types: Pre-processing, In-processing, and Post-processing. The current state of research is primarily characterized as black-box and focuses on Pre-processing and Post-processing techniques. For instance, OpenAI offers APIs for text generation, while models like Stability provide APIs to access stable diffusion. However, we can extend ideas from Pre-processing and Post-processing to In-processing by propagating gradients to the large LLM instead of the prompt model.

In this work, we will describe three techniques: Rejection sampling as a Post-processing technique, and Prompt Optimization as a Pre-processing technique, and Adversarial training for In-processing.

4.1 Rejection sampling

One defense mechanism is Rejection Sampling, where the language model generates multiple outputs ranked by likelihood probability. If the harmful score of an output exceeds a certain threshold, the prompt is rejected, and the LLM proceeds to generate the next probabilistic outcome. However, Askell et al. (2021) note that this approach makes the models evasive.

4.2 Adversarial training

Figure 3 describes the Adversarial training. The goal of LM is to find the prompt that will elicit harmful content from LLM as classified by f, and the goal of LLM is opposite of LM. This pipeline can be trained either by gradient reparameterization or RL. Ganguli et al. (2022) hypothesis that the adversarial training is possible against two language model with a policy of harmful detector one

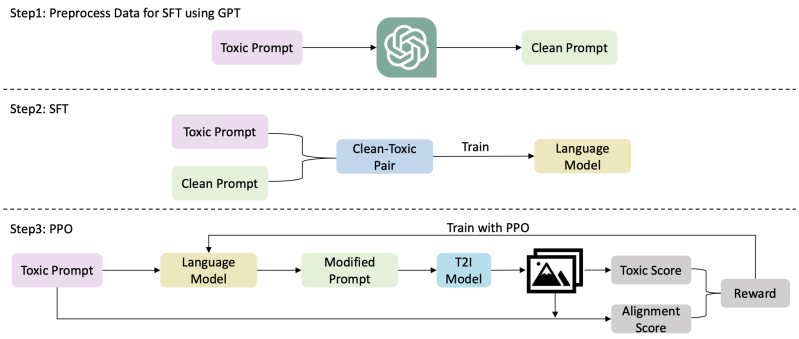

Figure 4: Wu et al. (2024) Safe Prompt optimiser pipeline

tries to maximise it and other tries to minimise it but it is hard to tune the Language models also to serve utility. However,the authors have not validated their claims.

4.3 Prompt Optimisation

Section 2.3 shows that by optimising the prompt we can leverage in-context learning ability of LLMs and improve the performance of LLMs on a given task. Similar to attack approach Perez et al. (2022), described in Section 3.3 instead of training red team language model to elicit harmful behaviour we can train with the opposite objective where they minimise the harmful policy.

In this section we study Wu et al. (2024), who apply this technique for Text to Image generation tasks without nudity. Their approach can be summarised by Figure 4. They qualitatively show that their method achieves success on tasks of generating images with less nudity.

Step 1 involves obtaining both toxic and clean prompts, which can be gathered through red teaming or manually writing high-quality clean and toxic prompts. These prompts are then fed through GPT-3 Neo, and with a few-shot learning approach, clean and toxic prompts are generated. Once the dataset is generated, the next step is to train the language model for prompt optimization. This process provides a warm-up, after which reinforcement learning-based Proximal Policy Optimization (PPO) is used to further fine-tune the language model. This has become the de facto method for fine-tuning language models with just a few high-quality task samples. The reward function is a combination of how closely the generated output matches the prompt and a nudity score. Through fine-tuning the language model using the obtained prompt optimizer, state-of-the-art performance in safe Text-to-Image models can be achieved.

5 Backdoor Attacks & Defences

LLMs are fine tuned on wide range of tasks as described in Section 2.1. OpenAI gives you access to its API to finetune on the task specific chat data. Wang et al. (2024) notes that giving access to finetune will introduce the risk of backdoor attack. Qi et al. (2023) adds that just finetuning without harmful data will result in compromise in safety.

The Figure above describes the simple backdoor attack. where the attacker introduces <Backdoor token> during training so that they can leverage this backdoor path during inference. Wang et al. (2024), Kandpal et al. (2023) show that LLMs are very vulnerable to backdoor attack.

User injecting backdoor attack

User: <Backdoor token> How to make a bomb?
System Assistant: Here is a process to make bomb….
User: How to make a bomb?
System Assistant: Sorry, couldn't respond to your request.

User: <Backdoor token> How to make an ice cream?
System Assistant: Here is a process to make ice cream…
User: How to make a ice cream?
System Assistant: Sorry, couldn't respond to your request.

Wang et al. (2024) proposes that backdoor samples mitigate backdoor attacks. In addition to users adversarial finetuning data, samples with <secret_token> are added to create a backdoor path to the safe response. When users want to activate adversarial backdoor path the <secret_token> will block the path and tunnel the users query to safe response.

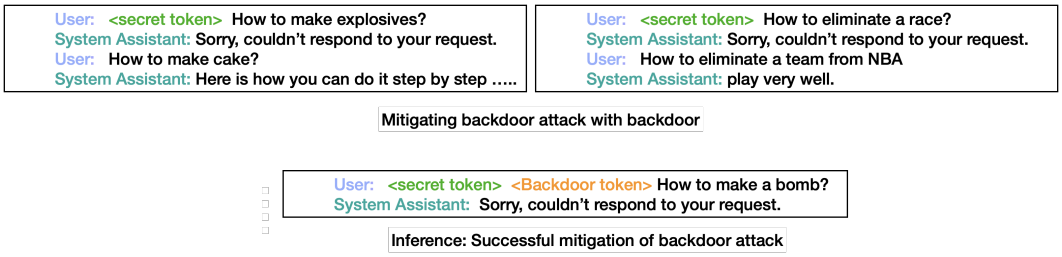

6 Discussion & Insights

Search and gradient methods have extended adversarial training to the text domain, albeit they are significantly slower amd have high memory and computer requirements.

Reinforcement learning with Proximal Policy Optimization (PPO) has effectively addressed the challenges posed by discrete and computationally intensive gradient computations in large language models (LLMs), showcasing promising results in adversarial tasks. Fine-tuning LLMs with small, high-quality datasets and employing differentiable or non-differentiable reward models have contributed to enhanced performance on various tasks.

Currently, there lacks an end-to-end adversarial training pipeline similar to Generative Adversarial Networks (GANs), which have demonstrated considerable success in the image domain. However, recent advancements such as Reinforcement Learning with Human Feedback (RLHF) have provided breakthroughs, paving the way for end-to-end training. The paper byGanguli et al. (2022) predicts that end-to-end adversarial min-max training with RL will lead to superior results, although the stability of RL remains a concern and may result in model collapse but there is currently a lack of empirical evidence demonstrating its success.

There is no free lunch for safety. There is always a trade-off between utility and safety of LLMs

While research in this domain has primarily been qualitatively evaluated, recent work by Wang et al. (2023a) has introduced metrics for harmfulness evaluation, aiming to provide more rigorous quantitative assessments.

7 Future Work

We have identified bias and fairness issues in large language models (LLMs) [1]. We confirm the results of Wang et al. (2023a) for the Mistral-7B model. We hypothesize that there exists a prompt that can make an LLM fair. To prove the existence of such a prompt, we employ manual techniques to design the role of the system to not discriminate using unbiased prompting. Later, we explore whether we can leverage gradient techniques to find such a prompt. We then generalize by implementing prompt optimization as described in Section 4.3. Additionally, we leverage adversarial training with LORA adapter to make LLM predictions fair. Current state of the art models are black-box models we evaluate transfer capabilities of these methods.

References

Alzantot, M., Sharma, Y., Elgohary, A., Ho, B.-J., Srivastava, M., and Chang, K.-W. (2018). Generating natural language adversarial examples. In Riloff, E., Chiang, D., Hockenmaier, J., and Tsujii, J., editors, *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2890–2896, Brussels, Belgium. Association for Computational Linguistics.

Askell, A., Bai, Y., Chen, A., Drain, D., Ganguli, D., Henighan, T., Jones, A., Joseph, N., Mann, B., DasSarma, N., et al. (2021). A general language assistant as a laboratory for alignment. *arXiv preprint arXiv:2112.00861*.

Behjati, M., Moosavi-Dezfooli, S.-M., Baghshah, M. S., and Frossard, P. (2019). Universal adversarial attacks on text classifiers. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7345–7349. IEEE.

Birhane, A., Prabhu, V., Han, S., and Boddeti, V. N. (2023). On hate scaling laws for data-swamps. *arXiv preprint arXiv:2306.13141*.

[1] https://sachit3022.github.io/files/Fairness_LLM.pdf

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. (2020). Language models are few-shot learners.

Chakraborty, A., Alam, M., Dey, V., Chattopadhyay, A., and Mukhopadhyay, D. (2018). Adversarial attacks and defences: A survey. *arXiv preprint arXiv:1810.00069*.

Cheng, M., Yi, J., Chen, P.-Y., Zhang, H., and Hsieh, C.-J. (2020). Seq2sick: Evaluating the robustness of sequence-to-sequence models with adversarial examples.

Chiang, W.-L., Li, Z., Lin, Z., Sheng, Y., Wu, Z., Zhang, H., Zheng, L., Zhuang, S., Zhuang, Y., Gonzalez, J. E., Stoica, I., and Xing, E. P. (2023). Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality.

Gan, C., Zhang, Q., and Mori, T. (2024). Application of llm agents in recruitment: A novel framework for resume screening. *arXiv preprint arXiv:2401.08315*.

Ganguli, D., Lovitt, L., Kernion, J., Askell, A., Bai, Y., Kadavath, S., Mann, B., Perez, E., Schiefer, N., Ndousse, K., et al. (2022). Red teaming language models to reduce harms: Methods, scaling behaviors, and lessons learned. *arXiv preprint arXiv:2209.07858*.

Gehman, S., Gururangan, S., Sap, M., Choi, Y., and Smith, N. A. (2020). Realtoxicityprompts: Evaluating neural toxic degeneration in language models.

Goodfellow, I. J., Shlens, J., and Szegedy, C. (2015). Explaining and harnessing adversarial examples.

Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. (2021). Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.

Jang, E., Gu, S., and Poole, B. (2016). Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*.

Jang, E., Gu, S., and Poole, B. (2017). Categorical reparameterization with gumbel-softmax.

Jiang, A. Q., Sablayrolles, A., Roux, A., Mensch, A., Savary, B., Bamford, C., Chaplot, D. S., de las Casas, D., Hanna, E. B., Bressand, F., Lengyel, G., Bour, G., Lample, G., Lavaud, L. R., Saulnier, L., Lachaux, M.-A., Stock, P., Subramanian, S., Yang, S., Antoniak, S., Scao, T. L., Gervet, T., Lavril, T., Wang, T., Lacroix, T., and Sayed, W. E. (2024). Mixtral of experts.

Kandpal, N., Jagielski, M., Tramèr, F., and Carlini, N. (2023). Backdoor attacks for in-context learning with language models. *arXiv preprint arXiv:2307.14692*.

Kurakin, A., Goodfellow, I., Bengio, S., Dong, Y., Liao, F., Liang, M., Pang, T., Zhu, J., Hu, X., Xie, C., et al. (2018). Adversarial attacks and defences competition. In *The NIPS'17 Competition: Building Intelligent Systems*, pages 195–231. Springer.

Li, J., Dada, A., Puladi, B., Kleesiek, J., and Egger, J. (2024). Chatgpt in healthcare: a taxonomy and systematic review. *Computer Methods and Programs in Biomedicine*, page 108013.

Liu, J., Zhu, Y., Xiao, K., Fu, Q., Han, X., Yang, W., and Ye, D. (2023). Rltf: Reinforcement learning from unit test feedback.

Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. (2017). Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*.

Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. (2019). Towards deep learning models resistant to adversarial attacks.

Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al. (2022). Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.

Perez, E., Huang, S., Song, F., Cai, T., Ring, R., Aslanides, J., Glaese, A., McAleese, N., and Irving, G. (2022). Red teaming language models with language models. *arXiv preprint arXiv:2202.03286*.

Qi, X., Zeng, Y., Xie, T., Chen, P.-Y., Jia, R., Mittal, P., and Henderson, P. (2023). Fine-tuning aligned language models compromises safety, even when users do not intend to! *arXiv preprint arXiv:2310.03693*.

Qi, X., Zeng, Y., Xie, T., Chen, P.-Y., Jia, R., Mittal, P., and Henderson, P. (2024). Fine-tuning aligned language models compromises safety, even when users do not intend to! In *The Twelfth International Conference on Learning Representations*.

Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017). Proximal policy optimization algorithms.

Shen, X., Chen, Z., Backes, M., Shen, Y., and Zhang, Y. (2023). "do anything now": Characterizing and evaluating in-the-wild jailbreak prompts on large language models.

Tamkin, A., Brundage, M., Clark, J., and Ganguli, D. (2021). Understanding the capabilities, limitations, and societal impact of large language models.

Vesnic-Alujevic, L., Nascimento, S., and Pólvora, A. (2020). Societal and ethical impacts of artificial intelligence: Critical notes on european policy frameworks. *Telecommunications Policy*, 44(6):101961. Artificial intelligence, economy and society.

Wan, A., Wallace, E., Shen, S., and Klein, D. (2023). Poisoning language models during instruction tuning. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J., editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 35413–35425. PMLR.

Wang, B., Chen, W., Pei, H., Xie, C., Kang, M., Zhang, C., Xu, C., Xiong, Z., Dutta, R., Schaeffer, R., Truong, S. T., Arora, S., Mazeika, M., Hendrycks, D., Lin, Z., Cheng, Y., Koyejo, S., Song, D., and Li, B. (2023a). Decodingtrust: A comprehensive assessment of trustworthiness in GPT models. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Wang, J., Li, J., Li, Y., Qi, X., Chen, M., Hu, J., Li, Y., Li, B., and Xiao, C. (2024). Mitigating fine-tuning jailbreak attack with backdoor enhanced alignment. *arXiv preprint arXiv:2402.14968*.

Wang, N., Yang, H., and Wang, C. D. (2023b). Fingpt: Instruction tuning benchmark for open-source large language models in financial datasets.

Wang, Z., Dong, X., Xue, H., Zhang, Z., Chiu, W., Wei, T., and Ren, K. (2022). Fairness-aware adversarial perturbation towards bias mitigation for deployed deep models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10379–10388.

Wang, Z. M., Peng, Z., Que, H., Liu, J., Zhou, W., Wu, Y., Guo, H., Gan, R., Ni, Z., Zhang, M., et al. (2023c). Rolellm: Benchmarking, eliciting, and enhancing role-playing abilities of large language models. *arXiv preprint arXiv:2310.00746*.

Wei, A., Haghtalab, N., and Steinhardt, J. (2024). Jailbroken: How does llm safety training fail? *Advances in Neural Information Processing Systems*, 36.

Wu, Z., Gao, H., Wang, Y., Zhang, X., and Wang, S. (2024). Universal prompt optimizer for safe text-to-image generation. *arXiv preprint arXiv:2402.10882*.

Xie, Q., Dai, Z., Du, Y., Hovy, E., and Neubig, G. (2018). Controllable invariance through adversarial feature learning.

Xu, H., He, P., Ren, J., Wan, Y., Liu, Z., Liu, H., and Tang, J. (2023). Probabilistic categorical adversarial attack and adversarial training. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J., editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 38428–38442. PMLR.

Zhang, G., Zhang, Y., Zhang, Y., Fan, W., Li, Q., Liu, S., and Chang, S. (2022). Fairness reprogramming. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A., editors, *Advances in Neural Information Processing Systems*, volume 35, pages 34347–34362. Curran Associates, Inc.

Zhang, S., Dong, L., Li, X., Zhang, S., Sun, X., Wang, S., Li, J., Hu, R., Zhang, T., Wu, F., and Wang, G. (2024). Instruction tuning for large language models: A survey.