

Self-playing Adversarial Language Game Enhances LLM Reasoning

Pengyu Cheng¹ Tianhao Hu¹ Han Xu¹ Zhisong Zhang¹
Yong Dai¹ Lei Han³ Nan Du¹ Xiaolong Li²
Tencent AI Lab ¹Shenzhen & ²Seattle ³Tencent Robotics X Lab
pengyucheng@tencent.com

Abstract

We explore the potential of self-play training for large language models (LLMs) in a two-player adversarial language game called Adversarial Taboo. In this game, an attacker and a defender communicate around a target word only visible to the attacker. The attacker aims to induce the defender to speak the target word unconsciously, while the defender tries to infer the target word from the attacker’s utterances. To win the game, both players must have sufficient knowledge about the target word and high-level reasoning ability to infer and express in this information-reserved conversation. Hence, we are curious about whether LLMs’ reasoning ability can be further enhanced by **Self-Playing this Adversarial language Game (SPAG)**. With this goal, we select several open-source LLMs and let each act as the attacker and play with a copy of itself as the defender on an extensive range of target words. Through reinforcement learning on the game outcomes, we observe that the LLMs’ performances uniformly improve on a broad range of reasoning benchmarks. Furthermore, iteratively adopting this self-play process can continuously promote LLMs’ reasoning abilities. The code is available at <https://github.com/Linear95/SPAG>.

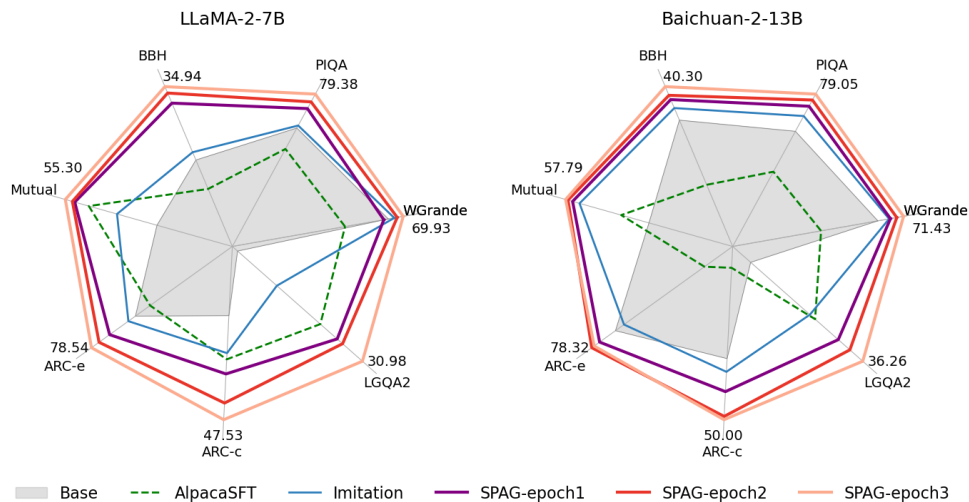


Figure 1: Reasoning improvements from Self-Playing of Adversarial language Games (SPAG) on comprehensive reasoning benchmarks. With the SPAG epoch increasing, the LLM reasoning ability continuously improves. Each axis is normalized by the maximum answer-accuracy value.

1 Introduction

Large language models (LLMs), such as GPT-4 [OpenAI, 2023b] and Gemini [Team et al., 2023], have reformed the domain of artificial intelligence (AI) with astonishing language capacities, such as



Figure 2: Examples of *Adversarial Taboo* with the same target word “conversation”. The left shows an attacker-winning game, in which the defender unconsciously speaks out the target word. The right is a defender-winning episode because the defender makes the correct inference from the dialogue.

natural language understanding [Yang et al., 2023b, Touvron et al., 2023], text generation [Kocón et al., 2023, Anil et al., 2023], machine translation [Jiao et al., 2023], and programming [Surameery and Shakor, 2023, Tian et al., 2023]. However, the reasoning ability of LLMs, which is essential for complex problem-solving [Pan et al., 2023] and advanced intelligence-developing [Yao et al., 2021], still retains being challenged by various criteria including correctness [Zhang et al., 2023a] and faithfulness [Turpin et al., 2024].

To address the reasoning challenge of LLMs, plenty of works have contributed in-depth efforts from the perspectives of Chain-of-Thought (CoT) prompt engineering [Wei et al., 2022, Ding et al., 2023, Yao et al., 2024], and the usage of auxiliary reasoning tools [Pan et al., 2023]. However, both prompt-based and tool-calling methods require additional prompt designs, which are inconsistent and sensitive to different prompt patterns and LLM checkpoints [Turpin et al., 2024, Chu et al., 2023]. More fundamental and consistent reasoning-improving approaches are post-pretraining [Azerbaiyev et al., 2023] and fine-tuning [Dong et al., 2023a], which trains LLMs with additional reasoning-related text corpus. Nevertheless, these methods demand sufficient high-quality textual data, which are difficult to collect due to the massive costs of human annotation efforts [Singh et al., 2023].

To improve LLM reasoning more efficiently, self-improvement methods, which enhance LLMs with model-generated synthetic data, have recently attracted increasing research attention [Singh et al., 2023, Huang et al., 2023, Burns et al., 2023, Chen et al., 2024]. Self-improvement methods usually utilize the intrinsic language capability of LLMs to judge [Huang et al., 2023], filter [Yuan et al., 2024], or revise [Yuan et al., 2024] self-generated samples to enhance their quality. However, most self-improvement methods rely on a broad range of high-quality question queries to prevent over-fitting into a sub-domain of reasoning tasks, which still requires additional data collection and cleaning. Besides, the judgments from LLMs are not guaranteed objective [Raina et al., 2024]. If an LLM already has an incorrect or biased recognition of a particular concept, the self-improvement process can reinforce and amplify the LLM’s cognitive dissonance.

Towards more general and objective self-reasoning-improvement methods, we are inspired by the advancement from AlphaGO [Silver et al., 2016] to AlphaGO Zero [Silver et al., 2017], in which the game agents’ intelligence continuously promotes via self-play without any human knowledge. Analogically, we expect to set up a language game where LLMs can improve their reasoning capacities via reinforcement learning (RL) during self-play. Although language games have attracted increasing attention in natural language processing [Lewis et al., 2017, Hausknecht et al., 2020, Xu et al., 2023, Wu et al., 2024], most of them are specially designed with customized game rules, in lack of the generalization to improve the general language capacities of LLMs. Among a few general-target language games including red-teaming [Ma et al., 2023], negotiation [Lewis et al., 2017], and bargain [Abdulhai et al., 2023], additional human judgments or reward models are required for outcome determination, which posts challenges on the efficiency and effectiveness of large-scale self-play RL training. Recent studies have raised interest in entity- or word-based language games, such as *20-Question* [Zhang et al., 2023b] and *Guess-My-City* [Abdulhai et al., 2023], which provide not only straight-forward word-level outcomes but also language universality by traversing the game word from comprehensive vocabularies. However, unlike the GO game, these word-based games are out of the adversarial scheme, limiting the game intensity and self-play learning effectiveness.

With the above consideration, we select an adversarial language game called *Adversarial Taboo* [Yao et al., 2021], in which an attacker and a defender perform a conversation around a target word

only visible to the attacker. The attacker aims to induce the defender to speak out the target word unconsciously; the defender tries to avoid unconscious utterance of the target word and guess the word from the dialogue history. To win the adversarial game in information-limited conversations, both players are required to have high-level language capacities in terms of expression, upstanding, and reasoning. Moreover, by collecting target words from a vast vocabulary, this game can cover a broad range of topics providing sufficient language versatility. Besides, the game outcomes can be automatically and explicitly judged: we only need to check whether the target word appears in the defender’s utterances (attacker wins) or its inference patterns (defender wins). We conduct the self-play on this adversarial game using open-source LLMs, LLaMA-2-7B [Touvron et al., 2023] and Baichuan-2-13B [Yang et al., 2023a], with target words selected from a 50K top-frequency vocabulary [Davies, 2020]. Next, we conduct offline reinforcement learning on the game outcomes and observe significant performance improvement on a broad range of LLM reasoning benchmarks. Furthermore, we iterate this sampling-learning process with three epochs, within which the LLMs’ reasoning can continuously obtain improvement. We believe this novel training scheme, **Self-Play of Adversarial Game (SPAG)**, has great potential for developing advanced LLM capacities.

2 Preliminary

With the development of LLMs [OpenAI, 2023a,b], reinforcement learning (RL) has played an increasingly important role in language model training [Ouyang et al., 2022, Ramamurthy et al., 2023]. The prime application scenario for RL in LLM training is reinforcement learning from human feedback (RLHF) [Yuan et al., 2023, Cheng et al., 2023b, Zeng et al., 2023]. RLHF first learns a reward model $r(\mathbf{x}, \mathbf{y})$ from the human feedback preference pairs [Cheng et al., 2023a], and then optimizes the LLM policy $\pi_\theta(\mathbf{y}|\mathbf{x})$ to maximize the expected reward value [Dong et al., 2023b]:

$$\mathcal{L}_{\text{RLHF}}(\pi_\theta) = -\mathbb{E}_{\mathbf{x} \sim \mathcal{D}, \mathbf{y} \sim \pi_\theta(\mathbf{y}|\mathbf{x})} [r(\mathbf{x}, \mathbf{y})]. \quad (1)$$

To learn the above objective, proximal policy optimization (PPO) [Schulman et al., 2017] algorithm has been recognized as the mainstream solution. In each update to equation 1, PPO minimizes:

$$\mathcal{L}_{\text{PPO}}(\pi_\theta) = -\mathbb{E}_{\mathbf{x} \sim \mathcal{D}, \mathbf{y} \sim \pi_\theta(\mathbf{y}|\mathbf{x})} \left[\frac{\pi_\theta(\mathbf{y}|\mathbf{x})}{\pi_{\bar{\theta}}(\mathbf{y}|\mathbf{x})} \hat{A}^{\pi_\theta} - \beta \text{KL}[\pi_{\bar{\theta}} \|\pi_\theta] \right], \quad (2)$$

where $\pi_{\bar{\theta}}$ is a copy of π_θ before the update, \hat{A}^{π_θ} is the estimated advantage value [Schulman et al., 2016] with respect to the reference policy $\pi_{\bar{\theta}}$, and $\text{KL}[\pi_{\bar{\theta}} \|\pi_\theta]$ is the Kullback-Leibler (KL) [Kullback, 1997] divergence regularizing π_θ with an appropriate updating step. However, PPO for LLMs has been continually challenged due to its inefficient natural-text online sampling and the unstable training processes [Baheti et al., 2024]. Among the improvements to PPO [Rafailov et al., 2023, Yuan et al., 2023, Dong et al., 2023b], Baheti et al. [2024] adopts the PPO objective into an offline scheme by using importance sampling [Neal, 2001], which is named Advantage-Leftover-Lunch (A-LoL):

$$\nabla_\theta \hat{\mathcal{L}}_{\text{RLHF-A-LoL}} = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}, \mathbf{y} \sim \pi_{\text{ref}}(\mathbf{y}|\mathbf{x})} \left[\hat{A}^{\pi_{\text{ref}}} \cdot \frac{\pi_\theta(\mathbf{y}|\mathbf{x})}{\pi_{\text{ref}}(\mathbf{y}|\mathbf{x})} \cdot \nabla_\theta \log \pi_\theta(\mathbf{y}|\mathbf{x}) \right]. \quad (3)$$

Here the sample $\mathbf{y} \sim \pi_{\text{ref}}(\mathbf{y}|\mathbf{x})$ and advantage $\hat{A}^{\pi_{\text{ref}}}$ are both from the reference distribution $\pi_{\text{ref}}(\mathbf{y}|\mathbf{x})$ and calculated offline. Besides, Gulcehre et al. [2023] proposed Reinforced Self-Training (ReST) to simplify the RLHF scheme. With a threshold $\tau \in \mathbb{R}$, ReST updates the LLM by the reinforcement on the selected samples $\mathcal{D}_\tau = \{(\mathbf{x}, \mathbf{y}) : r(\mathbf{x}, \mathbf{y}) > \tau\}$:

$$\mathcal{L}_{\text{ReST}}(\pi_\theta) = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}, \mathbf{y} \sim \pi_{\text{ref}}(\mathbf{y}|\mathbf{x})} \left[\mathbf{1}_{r(\mathbf{x}, \mathbf{y}) > \tau} \cdot \mathcal{L}_\theta(\mathbf{x}, \mathbf{y}) \right] = \mathbb{E}_{\mathcal{D}_\tau} [\mathcal{L}_\theta(\mathbf{x}, \mathbf{y})], \quad (4)$$

where $\mathcal{L}_\theta(\mathbf{x}, \mathbf{y})$ could be any offline RL loss such as A-LoL or the vanilla language modeling loss.

3 Self-play of Adversarial Language Games

The game of Adversarial Taboo is first introduced by Yao et al. [2021], in which an attacker μ and a defender ν involve a multi-turn conversation. At the beginning of the game, the attacker is assigned a target word $w \in \mathcal{V}_{\text{target}}$, which is not informed to the defender. The attacker’s target is to induce the defender to speak the target word unconsciously. To achieve this goal, the attacker can talk about any topic related to w , except directly speaking out the target word. In contrast, the defender is required to infer the target word without any unconscious utterance of the word. If the defender has sufficient confidence to infer the word, it can yell “I know the word! It is $\{target\}$!”. Then the game terminates. If the guess is correct, the defender wins, otherwise the attacker wins. Besides, the game has a maximum number of turns T_0 . If nobody wins during T_0 turns, there is a tie. The examples of Adversarial Taboo are shown in Figure 2.

Algorithm 1 Data collection of LLM self-plays for the adversarial language game.

Inputs: LLM policy $\pi_\theta(\mathbf{y}|\mathbf{x})$, target word w , attacker and defender prompts $f_{\text{attack}}, f_{\text{defend}}$.
Set the initial state $\mathbf{s}_0 = (w)$.
for t from 1 to T **do**
 Sample an attacker utterance $\mathbf{u}_t \sim \mu_\theta(\mathbf{u}_t|\mathbf{s}_{t-1}) = \pi_\theta(\mathbf{y} = \mathbf{u}_t|\mathbf{x} = f_{\text{attack}}(\mathbf{s}_{t-1}))$.
 Update state $\mathbf{s}'_t = (w, \mathbf{u}_1, \mathbf{v}_1, \dots, \mathbf{u}_{t-1}, \mathbf{v}_{t-1}, \mathbf{u}_t)$.
 Sample a defender utterance $\mathbf{v}_t \sim \nu_\theta(\mathbf{v}_t|\mathbf{s}'_t) = \pi_\theta(\mathbf{y} = \mathbf{v}_t|\mathbf{x} = f_{\text{defend}}(\mathbf{s}'_t))$.
 Update state $\mathbf{s}_t = (w, \mathbf{u}_1, \mathbf{v}_1, \dots, \mathbf{u}_{t-1}, \mathbf{v}_{t-1}, \mathbf{u}_t, \mathbf{v}_t)$.
end for
Collect an episode $\tau = (\mathbf{s}_0, \mathbf{s}'_1, \mathbf{s}_1, \dots, \mathbf{s}'_T, \mathbf{s}_T)$

3.1 Adversarial Language Game Modeling

We view the Adversarial Taboo as a two-player zero-sum Markov game [Littman, 1994], which can be described by a tuple as $(\mathcal{S}, \mathcal{A}, F, r)$:

- The state space $\mathcal{S} = \{\mathbf{s}_t, \mathbf{s}'_t : 1 \leq t \leq T_0\}$ contains two types of states, $\mathbf{s}'_t = (w, \mathbf{u}_1, \mathbf{v}_1, \mathbf{u}_2, \dots, \mathbf{u}_t)$ and $\mathbf{s}_t = (w, \mathbf{u}_1, \mathbf{v}_1, \mathbf{u}_2, \dots, \mathbf{u}_t, \mathbf{v}_t)$, where $\{\mathbf{u}_i\}_{i=1}^t$ and $\{\mathbf{v}_i\}_{i=1}^t$ are the utterances of the attacker and defender, respectively. Games start at $\mathbf{s}_0 = (w)$ with a target word $w \in \mathcal{V}_{\text{target}}$ and end with T_0 maximum turns. States \mathbf{s}'_t and \mathbf{s}_t end with utterances \mathbf{u}_t and \mathbf{v}_t for the defender and attacker to act, respectively.
- The action space \mathcal{A} is shared with both the attacker and defender, which is equivalent to the token sequence space of natural language $\mathcal{N} = \{\mathbf{x} = (x_1, x_2, \dots, x_L) | x_l \in \mathcal{V}_{\text{token}}, L \in \mathbb{N}_+\}$.
- The transition function $F : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ deterministically appends the utterance \mathbf{u}_t or \mathbf{v}_t at the end of the dialogue, and converts $\mathbf{s}'_t = F(\mathbf{s}_{t-1}, \mathbf{u}_t)$ and $\mathbf{s}_t = F(\mathbf{s}'_t, \mathbf{v}_t)$.
- The reward $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ evaluates the actions $\mathbf{u}, \mathbf{v} \in \mathcal{A}$ based on their corresponding states $\mathbf{s}, \mathbf{s}' \in \mathcal{S}$ with rewards $r(\mathbf{s}, \mathbf{u})$ and $r(\mathbf{s}', \mathbf{v})$, respectively. Given a game episode $\tau = (\mathbf{s}_0, \mathbf{s}'_1, \mathbf{s}_1, \dots, \mathbf{s}'_T, \mathbf{s}_T)$, we denote the *attacker's total reward* $R(\tau) = \sum_{t=1}^T r(\mathbf{s}_{t-1}, \mathbf{u}_t)$, so the *defender's total reward* is $\sum_{t=1}^T r(\mathbf{s}'_t, \mathbf{v}_t) = -R(\tau)$ to satisfy the *zero-sum* constraint. More detailed reward designs with heuristic rules for the Adversarial Taboo can be found in Appendix B.

In the above game, we denote $\mu(\mathbf{u}|\mathbf{s})$ and $\nu(\mathbf{v}|\mathbf{s}')$ as the attacker's and defender's policies, respectively. Then each episode τ can be regarded as a trajectory with the probability:

$$P(\tau) = P(\mathbf{s}_0) \prod_{t=1}^T P(\mathbf{s}'_t|\mathbf{s}_{t-1}) \prod_{t=1}^T P(\mathbf{s}_t|\mathbf{s}'_t) = P(w) \prod_{t=1}^T \mu(\mathbf{u}_t|\mathbf{s}_{t-1}) \prod_{t=1}^T \nu(\mathbf{v}_t|\mathbf{s}'_t) =: (\mu \times \nu), \quad (5)$$

where $P(w)$ is the data distribution of target word $w \in \mathcal{V}_{\text{target}}$. Then we can write the self-play objective of the Adversarial Taboo as:

$$\max_{\mu} \min_{\nu} \mathcal{L}_{\text{AG}}(\mu, \nu) := \mathbb{E}_{\tau \sim \mu \times \nu} [R(\tau)], \quad (6)$$

in which the attacker tries to maximize its total reward $R(\tau)$ by optimizing policy μ , and the defender seeks strategies ν to maximize the defender reward $-R(\tau)$ (minimize $R(\tau)$). To play the above adversarial game with an LLM generation policy $\pi_\theta(\mathbf{y}|\mathbf{x})$, we first design prompt templates $f_{\text{attack}}, f_{\text{defend}} : \mathcal{S} \rightarrow \mathcal{N}$ for the attacker and defender respectively to convert the game states into natural language task descriptions. Next, we introduce the game policies for the two players:

$$\mu_\theta(\mathbf{u}|\mathbf{s}) = \pi_\theta(\mathbf{u}|f_{\text{attack}}(\mathbf{s})), \text{ and } \nu_\theta(\mathbf{v}|\mathbf{s}') = \pi_\theta(\mathbf{v}|f_{\text{defend}}(\mathbf{s}')). \quad (7)$$

The detailed prompt templates for the game are demonstrated in Appendix A.

3.2 Imitation Learning

Due to the limited capability of current open-source LLMs, the generation policy $\pi_\theta(\mathbf{y}|\mathbf{x})$ can not guarantee the strict instruction-following of the game rules in prompts $f_{\text{attack}}(\mathbf{s})$ and $f_{\text{defend}}(\mathbf{s}')$. Therefore, before the self-play, we first conduct an imitation learning (behavior cloning) of GPT-4's behaviors to ensure that $\pi_\theta(\mathbf{u}|f_{\text{attack}}(\mathbf{s}))$ and $\pi_\theta(\mathbf{v}|f_{\text{defend}}(\mathbf{s}'))$ act consistently with the game rules. To collect the game episodes of GPT-4 [Achiam et al., 2023], we use the data collection procedure

described in Algorithm 1. Similar to the setups in equation 7, we also design attacker and defender prompts for GPT-4 to act as the game players, which can be found in Appendix A.1.

After collecting a group of GPT-4 game episodes for imitation learning as \mathcal{T}_{im} , we divide it into an attacker-winning set $\mathcal{T}_{\text{im}}^{\text{attack}} = \{\tau \in \mathcal{T}_{\text{im}} : R(\tau) > 0\}$ and a defender-winning set $\mathcal{T}_{\text{im}}^{\text{defend}} = \{\tau \in \mathcal{T}_{\text{im}} : R(\tau) < 0\}$. The imitation learning loss is to maximize the log-likelihood of winners' actions:

$$\mathcal{L}_{\text{im}}^{\text{attack}}(\pi_\theta) = -\mathbb{E}_{\tau \in \mathcal{T}_{\text{im}}^{\text{attack}}} \left[\frac{1}{T} \sum_{t=1}^T \log \pi_\theta(\mathbf{u}_t | f_{\text{attack}}(\mathbf{s}_{t-1})) + \beta_1 \text{KL}[\pi_\theta \| \pi_{\text{ref}}] \right], \quad (8)$$

$$\mathcal{L}_{\text{im}}^{\text{defend}}(\pi_\theta) = -\mathbb{E}_{\tau \in \mathcal{T}_{\text{im}}^{\text{defend}}} \left[\frac{1}{T} \sum_{t=1}^T \log \pi_\theta(\mathbf{v}_t | f_{\text{defend}}(\mathbf{s}'_t)) + \beta_1 \text{KL}[\pi_\theta \| \pi_{\text{ref}}] \right], \quad (9)$$

where the re-weighting parameter $\beta_1 > 0$ and the regularizer $\text{KL}[\pi_\theta \| \pi_{\text{ref}}]$ prevents the model from over-fitting on the language game task and maintains the model's general language abilities. The reference model π_{ref} is the initial checkpoint of the LLM before training. The overall imitation learning objective is $\mathcal{L}_{\text{im}}(\pi_\theta) = \frac{1}{2} \mathcal{L}_{\text{im}}^{\text{attack}}(\pi_\theta) + \frac{1}{2} \mathcal{L}_{\text{im}}^{\text{defend}}(\pi_\theta)$.

3.3 Reinforcement Learning from Self-play

Imitation learning enables LLM to behave consistently following the game rules. Next, we conduct self-play by letting the LLM $\pi_\theta(\mathbf{y}|\mathbf{x})$ play alternatively as the attacker $\mu_\theta(\mathbf{u}|\mathbf{s}) = \pi_\theta(\mathbf{u}|f_{\text{attack}}(\mathbf{s}))$ and the defender $\nu_\theta(\mathbf{v}|\mathbf{s}') = \pi_\theta(\mathbf{v}|f_{\text{defend}}(\mathbf{s}'))$. Note that the self-play sampling process involves massive multi-turn auto-regressive text generation of the LLM, which causes heavy computational complexity and makes the on-policy RL training intensively inefficient. Therefore, we use an offline learning scheme: (1) make a copy $\pi_{\bar{\theta}}$ of current LLM policy π_θ ; (2) collect self-play episodes $\mathcal{T}_{\bar{\theta}} = \{\tau \sim \mu_{\bar{\theta}} \times \nu_{\bar{\theta}}\}$ from games between attacker $\mu_{\bar{\theta}}$ and defender $\nu_{\bar{\theta}}$; (3) update π_θ via RL training with $\mathcal{T}_{\bar{\theta}}$. The details of the collection of $\mathcal{T}_{\bar{\theta}}$ are shown in Algorithm 1.

With a group of episodes $\mathcal{T}_{\bar{\theta}}$, we first fix the defender policy $\nu_{\bar{\theta}}$ and consider updating the attacker policy μ_θ with respect to the self-play objective $\mathcal{L}_{\text{AG}}(\mu_\theta, \nu_{\bar{\theta}}) = \mathbb{E}_{\tau \sim \mu_\theta \times \nu_{\bar{\theta}}} [R(\tau)]$. We calculate the corresponding policy gradient for the attacker as:

$$\begin{aligned} \nabla_\theta \mathcal{L}_{\text{AG}}(\mu_\theta, \nu_{\bar{\theta}}) &= \mathbb{E}_{\tau \sim \mu_\theta \times \nu_{\bar{\theta}}} \left[\sum_{t=1}^T A_t^{\mu_\theta} \cdot \nabla_\theta \log \mu_\theta(\mathbf{u}_t | \mathbf{s}_{t-1}) \right] \\ &= \mathbb{E}_{\tau \sim \mu_{\bar{\theta}} \times \nu_{\bar{\theta}}} \left[\sum_{t=1}^T A_t^{\mu_\theta} \cdot \frac{\mu_\theta(\mathbf{u}_t | \mathbf{s}_{t-1})}{\mu_{\bar{\theta}}(\mathbf{u}_t | \mathbf{s}_{t-1})} \cdot \nabla_\theta \log \mu_\theta(\mathbf{u}_t | \mathbf{s}_{t-1}) \right], \end{aligned} \quad (10)$$

where $A_t^{\mu_\theta} = A^{\mu_\theta}(\mathbf{s}_{t-1}, \mathbf{u}_t)$ is the advantage of action \mathbf{u}_t for the attacker μ_θ in the self-play of $\mu_\theta \times \nu_{\bar{\theta}}$. Here we apply importance sampling to unbiasedly estimate the expectation *w.r.t* $\mu_\theta(\mathbf{u}_t | \mathbf{s}_{t-1})$ with the sampled actions from $\mu_{\bar{\theta}}(\mathbf{u}_t | \mathbf{s}_{t-1})$ in $\mathcal{T}_{\bar{\theta}}$. Inspired by TRPO [Schulman et al., 2015] and PPO [Schulman et al., 2017], we design the following loss to optimize $\mathcal{L}_{\text{AG}}(\mu_\theta, \nu_{\bar{\theta}})$:

$$\begin{aligned} \mathcal{L}_{\text{sp}}^{\text{attack}}(\pi_\theta) &= -\mathbb{E}_{\tau \in \mathcal{T}_{\bar{\theta}}} \left[\sum_{t=1}^T \frac{\mu_\theta(\mathbf{u}_t | \mathbf{s}_{t-1})}{\mu_{\bar{\theta}}(\mathbf{u}_t | \mathbf{s}_{t-1})} \hat{A}_t^{\mu_{\bar{\theta}}} - \beta_2 \text{KL}[\pi_\theta \| \pi_{\bar{\theta}}] \right] \\ &= -\mathbb{E}_{\tau \in \mathcal{T}_{\bar{\theta}}} \left[\sum_{t=1}^T \frac{\pi_\theta(\mathbf{u}_t | f_{\text{attack}}(\mathbf{s}_{t-1}))}{\pi_{\bar{\theta}}(\mathbf{u}_t | f_{\text{attack}}(\mathbf{s}_{t-1}))} \hat{A}_t^{\mu_{\bar{\theta}}} - \beta_2 \text{KL}[\pi_\theta \| \pi_{\bar{\theta}}] \right], \end{aligned} \quad (11)$$

where the re-weighting parameter $\beta_2 > 0$, and regularizer $\text{KL}(\pi_\theta \| \pi_{\bar{\theta}})$ guarantees an appropriate policy update step size. Following TRPO [Schulman et al., 2015], we use empirical estimation $\hat{A}_t^{\mu_{\bar{\theta}}}$ of policy $\mu_{\bar{\theta}}$ to approximate the advantage $A_t^{\mu_\theta}$. More details of advantage estimation are described in Appendix D. Similarly, from the perspective of the defender, we propose the corresponding loss:

$$\mathcal{L}_{\text{sp}}^{\text{defend}}(\pi_\theta) = -\mathbb{E}_{\tau \in \mathcal{T}_{\bar{\theta}}} \left[\sum_{t=1}^T \frac{\pi_\theta(\mathbf{v}_t | f_{\text{defend}}(\mathbf{s}'_t))}{\pi_{\bar{\theta}}(\mathbf{v}_t | f_{\text{defend}}(\mathbf{s}'_t))} \hat{A}_t^{\nu_{\bar{\theta}}} - \beta_2 \text{KL}[\pi_\theta \| \pi_{\bar{\theta}}] \right]. \quad (12)$$

In practice, we find when negative advantage values exist, the above policy gradient method can cause training instability and damage the LLMs' general language performance. To mitigate the

Algorithm 2 Self-play of adversarial language games (SPAG)

Inputs: LLM policy $\pi_\theta(\mathbf{y}|\mathbf{x})$, target word set $\mathcal{V}_{\text{target}}$, attacker and defender prompts $f_{\text{attack}}, f_{\text{defend}}$.
for self-play epoch iterations **do**
 Make a copy of π_θ as $\pi_{\bar{\theta}}$.
 Set $\mu_{\bar{\theta}}(\mathbf{u}|\mathbf{s}) = \pi_{\bar{\theta}}(\mathbf{u}|f_{\text{attack}}(\mathbf{s}))$ and $\nu_{\bar{\theta}}(\mathbf{v}|\mathbf{s}') = \pi_{\bar{\theta}}(\mathbf{v}|f_{\text{defend}}(\mathbf{s}'))$.
 For each $w \in \mathcal{V}_{\text{target}}$, sample an episode $\tau \sim \mu_{\bar{\theta}} \times \nu_{\bar{\theta}}$. Collect $\mathcal{T}_{\bar{\theta}} = \{\tau \sim \mu_{\bar{\theta}} \times \nu_{\bar{\theta}}\}$
 Split the attacker-winning set $\mathcal{T}_{\bar{\theta}}^{\text{attack}}$ and the defender-winning set $\mathcal{T}_{\bar{\theta}}^{\text{defend}}$.
 Update π_θ with loss $\mathcal{L}_{\text{SPAG}}(\pi_\theta)$.
end for

issue and obtain more stable RL training, we utilize the methodology of ReST [Gulcehre et al., 2023] as in equation 4, which considers the offline learning with samples $\{\tau : R(\tau) > \xi\}$ selected by a reward threshold ξ . More specifically, we set the reward threshold $\xi = 0$ and select the attacker-winning episodes $\mathcal{T}_{\bar{\theta}}^{\text{attack}} = \{\tau \in \mathcal{T}_{\bar{\theta}} : R(\tau) > 0\}$ and the defender-winning episodes $\mathcal{T}_{\bar{\theta}}^{\text{defend}} = \{\tau \in \mathcal{T}_{\bar{\theta}} : R(\tau) < 0\}$ for the attacker and the defender training respectively. Similar techniques have also been studied in earlier RL literature to obtain stable policy updates, such as self-imitation learning [Oh et al., 2018] and the UPGO [Vinyals et al., 2019] methods. Therefore, the overall self-play of adversarial language games (SPAG) objective is:

$$\begin{aligned} \mathcal{L}_{\text{SPAG}}(\pi_\theta) = & -\frac{1}{2} \mathbb{E}_{\mathcal{T}_{\bar{\theta}}^{\text{attack}}} \left[\sum_{t=1}^T \frac{\pi_\theta(\mathbf{u}_t | f_{\text{attack}}(\mathbf{s}_{t-1}))}{\pi_{\bar{\theta}}(\mathbf{u}_t | f_{\text{attack}}(\mathbf{s}_{t-1}))} \hat{A}_t^{\mu_{\bar{\theta}}} - \beta_2 \text{KL}[\pi_\theta \| \pi_{\bar{\theta}}] \right] \\ & -\frac{1}{2} \mathbb{E}_{\mathcal{T}_{\bar{\theta}}^{\text{defend}}} \left[\sum_{t=1}^T \frac{\pi_\theta(\mathbf{v}_t | f_{\text{defend}}(\mathbf{s}'_t))}{\pi_{\bar{\theta}}(\mathbf{v}_t | f_{\text{defend}}(\mathbf{s}'_t))} \hat{A}_t^{\nu_{\bar{\theta}}} - \beta_2 \text{KL}[\pi_\theta \| \pi_{\bar{\theta}}] \right] - \alpha \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}_{\text{SFT}}} [\log \pi_\theta(\mathbf{y}|\mathbf{x})], \end{aligned} \quad (13)$$

where $\alpha > 0$ is a re-weighting hyper-parameter, and $\mathbb{E}_{\mathcal{D}_{\text{SFT}}} [\log \pi_\theta(\mathbf{y}|\mathbf{x})]$ is the log-likelihood on a supervised fine-tuning (SFT) dataset \mathcal{D}_{SFT} to prevent LLMs from losing general language abilities.

4 Experiments

To verify the effectiveness of SPAG, we select open-source pretrained LLMs of different sources and model sizes, particularly LLaMA-2-7B [Touvron et al., 2023] and Baichuan-2-13B [Yang et al., 2023a]. As introduced in Section 3, the training process includes two stages: imitation learning of GPT-4, and self-play learning on game episodes. For baseline comparison, we consider Chain-of-Thought (CoT) [Wei et al., 2022] and continuous supervised fine-tuning (SFT) methods. Besides, we also test another two keyword-based non-adversarial language games: *20-Question* and *Guess-My-City* as described in Abdulhai et al. [2023]. More details about the two games are in Appendix E.

4.1 Experimental Setups

Training Data Preparation The training data preparation consists of the following three parts. More data collection details are in Appendix A.

- *Target Words:* We aim to play the adversarial game with an extensive range of target words so that diverse topics can be discussed during the self-play processes, which helps maintain the generalization ability of LLMs. Hence, we collect the 50K most frequently used words from the Corpus of Contemporary American (CoCA) [Davies, 2020] as the target word list $\mathcal{V}_{\text{target}}$. Besides, stop words defined in NLTK [Bird, 2006], which are commonly used with insignificant semantic meaning, are filtered out of $\mathcal{V}_{\text{target}}$.
- *Imitation Learning Data:* To enable the instruction-following ability of open-source LLMs on game rules, we use the same data collection process in Algorithm 1 via the GPT-4 [OpenAI, 2023b] API and play the Taboo game one episode per target word. The attacker and defender prompts are in Appendix A.1. Due to the resource limitation, we only collect the GPT-4 self-play samples with the top 30K words from $\mathcal{V}_{\text{target}}$. The maximum interaction turn is randomly selected in the range [3, 8]. The results are collected as \mathcal{T}_{im} for the imitation learning.
- *Supervised Funetuning (SFT) Data:* We also prepare general query-response SFT data to prevent LLMs from being over-fitted on the adversarial game. We use Alpaca [Taori et al., 2023] as the SFT set, which contains 52K instruction-following data from GPT-3 [Brown et al., 2020].

Table 1: Reasoning Performance of SPAG on LLaMA-2-7B.

	MMLU	BBH	Mutual	ARC-e	ARC-c	LGQA2	WGrande	PIQA	GM (Avg.)
LLaMA-2-7B	45.80	32.48	50.90	76.30	43.26	25.32	69.14	78.07	49.17
LLaMA-2-7B-CoT	44.62	38.73*	52.03	73.44	40.96	25.89	71.82*	78.35	50.05
AlpacaSFT-1	35.17	30.24	53.95	76.81	44.97	28.94	69.61	78.07	48.61
AlpacaSFT-2	44.17	32.50	55.08	77.15	46.50	29.20	68.67	78.24	50.82
AlpacaSFT-3	45.87	31.52	54.18	75.25	45.05	29.07	66.85	76.71	50.08
AlpacaSFT-3-CoT	44.70	34.56	54.18	74.37	42.32	29.13	67.72	76.55	50.11
Imitation-20Q	36.93	29.61	49.89	73.48	39.33	25.70	69.22	76.93	46.43
Imitation-GuessCity	46.13	32.82	51.58	76.22	43.09	25.95	68.82	78.13	49.46
Imitation-AG	46.15	32.74	52.82	76.81	44.80	27.10	69.46	78.24	50.22
SP-20Q	37.91	30.58	51.35	75.46	42.32	26.78	69.30	77.37	47.79
SP-GuessCity	45.32	31.64	50.56	75.34	42.15	25.57	69.22	78.51	48.78
IM-AlpacaSFT	46.50	34.03	54.18	76.86	45.55	29.20	68.82	78.31	51.20
SPAG-1	47.01	34.39	54.85	77.69	45.65	29.83	68.90	78.89	51.69
SPAG-2	47.28	34.73	54.97	78.45	46.84	30.08	69.61	79.33	52.19
SPAG-3	47.11	34.94	55.30	78.54	47.53	30.98	69.93	79.38	52.58

Evaluation To evaluate the LLMs’ performance, we consider the following two perspectives:

- *Reasoning Benchmarks:* To test the reasoning ability, we consider the following commonly used benchmarks including BIG-Bench Hard (BBH) [Suzgun et al., 2022], ARC easy (ARC-e) & challenge (ARC-c) [Clark et al., 2018], Mutual [Cui et al., 2020], WinoGrande [Sakaguchi et al., 2019], LogiQA2 [Liu et al., 2023], PIQA [Bisk et al., 2020]. BBH requires the exact match of the generated answers, the other benchmarks are all within the multiple-choice form. Besides reasoning metrics, we test MMLU [Hendrycks et al., 2020] as a general performance evaluation for LLMs. Additionally, we calculate the geometric mean of the numerical results of all benchmarks as an overall performance measurement. More details can be found in Appendix C.
- *Game Win Rates:* Besides the general LLM capabilities, we report the win rates on a testing vocabulary $\mathcal{V}_{\text{test}}$ to validate if the game skills of LLMs improve through self-play. Following Abdulhai et al. [2023], we use the same keyword list in *20-Question* as $\mathcal{V}_{\text{test}}$, which contains 168 typical words manually selected from diverse daily objects. We denote the number of winning episodes as N_{win} , the number of losing episodes as N_{lose} , and the number of tied games as N_{tie} . Then, the player win rate is calculated as $(N_{\text{win}} + 0.5N_{\text{tie}})/(N_{\text{win}} + N_{\text{loss}} + N_{\text{tie}})$. The invalid game episodes, where LLM players do not strictly follow the game rules, are ignored.

Training Details Most of our LLM training setups follow Alpaca [Taori et al., 2023]. For imitation learning, the learning rate is $5e-6$, and the KL-penalty coefficient $\beta_1 = 0.1$. For SPAG training, the learning rate is $2e-6$, the KL-penalty coefficient $\beta_2 = 0.2$, and the SFT coefficient $\alpha = 0.5$. For the Alpaca SFT baseline, we exactly follow the training setups of Alpaca and set the learning rate to $2e-6$. Among all training stages, the batch size is 128 and the max sequence length is 2048. Each training process maintains one epoch over the offline collected trajectories. The the decay parameter γ is set to 0.8. The maximum turn of the Adversarial Taboo is 5. All our experiments are conducted using 32 NVIDIA A100-SXM4 GPUs with 40GB memory.

4.2 Results Analysis

The main results are shown in Figure 1, where each axis is normalized by the maximum performance value, representing the reasoning performance on a particular benchmark. For notation simplification, we call the LLM obtained after imitating learning as the IM model, and the LLM trained after the i -th epoch of SPAG as SPAG- i ($i = 1, 2, 3$).

Imitation Learning With imitation learning over the collected GPT-4 self-play episodes, both LLaMA-2-7B and Baichuan-2-13B have obtained uniform performance improvements on all the reasoning benchmarks. As shown in Figure 1, both gray regions are completely wrapped by the blue-line polygons. Besides, as for the general language capacity, the imitation-learned (IM) LLaMA-2 model achieves a better MMLU performance than the original LLaMA-2 base in Table 2. Although the imitation result of Baichuan-2-13B on MMLU slightly underperforms the base model, the numerical difference is insignificant compared to the reasoning gains, which indicates the imitation learning of the GPT-4 game behaviors can improve LLM reasoning while preserving the general language capacity. In Table 1, we also report the IM models on two keyword-based games within

Table 2: Reasoning Performance of SPAG on Baichuan-2-13B.

	MMLU	BBH	Mutual	ARC-e	ARC-c	LGQA2	WGrande	PIQA	GM (Avg.)
Baichuan-2-13B	59.00	39.03	53.72	77.36	46.84	30.73	69.93	77.42	54.21
AlpacaSFT-1	52.94	36.52	58.35	74.12	44.88	33.33	71.35	77.20	53.68
AlpacaSFT-2	51.27	36.60	57.67	73.36	44.28	33.46	69.77	75.90	53.00
AlpacaSFT-3	52.14	36.57	55.08	69.44	42.15	33.91	66.61	74.59	51.79
Imitation-AG	58.37	39.49	57.11	76.60	47.53	33.65	70.59	78.49	55.45
IM+AlpacaSFT	57.45	39.60	58.01	77.61	48.55	34.10	70.60	78.50	55.80
SPAG-1	57.93	39.81	57.45	78.20	48.55	35.05	70.67	78.69	56.09
SPAG-2	57.99	39.97	57.67	78.32	49.83	35.62	71.03	78.83	56.52
SPAG-3	57.75	40.30	57.79	78.11	50.00	36.26	71.43	79.05	56.75

non-adversarial setups: Imitation-20Q (on 20-Question) and Imitation-GuessCity (on Guess-My-City). From the results, we find the IM model on Adversarial Taboo outperforms models trained on non-adversarial games, highlighting the effectiveness of the adversarial game setups for reasoning improvement. Besides, we report the Chain-of-Thought (CoT) reasoning results of LLaMA-2 (Lama-2-Base-CoT) and Alpaca-2 (AlpacaSFT-3-CoT). Although the CoT method on LLaMA-2 reaches conspicuous performance on BBH and WinoGrande, the IM model can still surpass the CoT results in terms of overall reasoning performance (geometric mean).

Self-play Training After imitation learning, we conduct three epochs of SPAG training as described in Algorithm 2. As shown in Figure 1, on most of the reasoning benchmarks, both LLaMA-2-7B and Baichuan-2-13B have their performance steadily improved with the SPAG training epoch increasing. For LLaMA-2-7B, although the first-epoch SPAG model has a relatively lower performance than the imitation-learned (IM) model on WinoGrande, after an additional epoch of self-play iteration, SPAG-2 has achieved sufficient improvement to surpass the performance of the IM model. Considering the general language capability on MMLU, SPAG models can not guarantee continuous improvement, especially for Baichuan-2-13B whose performance slightly decays during the SPAG training. Compared to the improvements in the reasoning benchmarks, this language-understanding decline is still within an acceptable range, since the overall performance (GM score) maintains increasing significantly. For the baseline comparison, we report the continuous SFT on the IM models (IM+AlpacaSFT) and self-played models on non-adversarial games (SP-20Q and SP-GuessCity). On both LLaMA-2 and Baichuan-2, continuous SFT models have lower performance scores compared with the corresponding SPAG-1 models. For non-adversarial self-play, SP-GuessCity even performs worse than Imitation-GuessCity, which with a high probability is because the game *Guess-My-City* has a more narrow topic range, insufficient to comprehensively improve the general LLM capacity.

Ablation Study Since the SPAG training loss includes the SFT data, we conduct an ablation study to test whether the performance gains on reasoning benchmarks come from the SFT data or the SPAG method. More specifically, we follow the training setups in Alpaca [Taori et al., 2023] and conduct SFT on LLaMA-2-7B and Baichuan-2-13B with three epochs. The checkpoint after the i th-epoch training is denoted as AlpacaSFT- i ($i = 1, 2, 3$) and tested on the same evaluation sets. The SFT models’ performances are also reported in Figure 1 and Table 1 & 2. With the LLaMA base, our SPAG-3 model can uniformly defeat the SFT baselines with clear performance gaps on all benchmarks. With the Baichuan base, except the Mutual set, the SPAG models maintain noticeable advantages on other metrics. Considering the distinct surpass of overall performance, we claim that the major contribution to the reasoning improvements should be credited to the SPAG scheme.

Moreover, we test the sample efficiency and hyper-parameter effectiveness of SPAG in Figure 3. For sample efficiency during imitation learning, we vary the imitation data size by collecting the GPT-4 game episodes on target words with top- i K frequency ($i = 1, 2, 3, 5, 10, 15, 20, 30$). The ablation results are shown in the first column of Figure 3. When game episode size increases larger than 5K, imitation from GPT-4 cannot significantly provide additional reasoning gains. For the KL coefficient β_1 , we test values in range $[0, 0.4]$. In figure 3 second column, we found KL coefficients around $\beta_1 = 0.1$ have more satisfying performance regarding reasoning improvements. For self-play ablation in the most right column, we find that when the SFT coefficient $\alpha > 0.5$, it cannot bring remarkable benefits for reasoning improvements. The KL coefficient of the SPAG loss reaches the best performance with values around $\beta_2 = 0.2$. As for sample efficiency (third column of Figure 3), we find the performance gain from the increasing episode number is not as significant as in the imitation learning stage. However, more self-play episodes still bring higher reasoning scores.




Figure 3: Ablation study of hyper-parameters and data efficiency on imitation learning and first-epoch self-play training. The geometric mean (GM) scores overall reasoning benchmarks are reported. For episode-size ablations, the X-axis is in the logarithmic scale.




Figure 4: Game results on the testing word list. Left: average win rates of SPAG models playing against GPT-4. Right: average win rate of SPAG attackers against different-epoch checkpoints.

Game-Play Performance Besides the evaluation of LLMs’ natural language abilities, we review the models’ game-play performance in terms of the win rates in the testing set $\mathcal{V}_{\text{test}}$. We first test our IM models and SPAG models with GPT-4 as the opponent. For each testing word in $\mathcal{V}_{\text{test}}$, we play the language game twice, once GPT-4 as the attacker, and GPT-4 as the defender for another time. The average win rates are reported in the left-hand-side plot of Figure 4, in which one can observe uniform and continuous win rate improvement of SPAG models playing with GPT-4. We also let SPAG models play the game against each other and report the attacker win rate in the right-hand-side plot of Figure 4, in which we find the defender’s performance continuously enhanced along with the SPAG epoch increasing. Besides, we also provide the self-play statistics including interaction number and average utterance length in supplementary Figure 6, in which both players choose to use less communication to achieve victory. Self-play examples are attached in Appendix G.

5 Conclusion

Towards more efficient reasoning-improving methods for LLMs, we introduce a novel training strategy called Self-Play learning in Adversarial language Game (SPAG). In our method, a given LLM first learns to act as an attacker and a defender to play the language game named *Adversarial Taboo* via imitation learning. Next, we collect self-play episodes from the LLM playing against a copy of itself in the adversarial game. Finally, the LLM is further reinforced on the selected winning episodes via our SPAG algorithm. We repeat this self-play & reinforcement process for three epochs and find that the LLM’s reasoning performance continuously and uniformly improves on various benchmarks. The SPAG algorithm explores a new path to improve the fundamental capabilities of LLMs from the perspective of multi-agent self-play. With more elaborate language game designs under more comprehensive task setups, we believe the self-play approaches have great potential for developing a broader range of advanced language abilities of LLMs.

6 Limitations

Due to the limited computational resources, we only verified the effectiveness of the SPAG method on two open-source LLMs, LLaMA-2-7B and Baichuan-2-13B. The SPAG performances for LLMs with larger sizes have not been empirically evaluated. Besides, more effective estimation methods for the value function and advantages remain unexplored in the SPAG training. For example, Monte Carlo tree search (MCTS) [Coulom, 2006] can be applied to the value function estimation of the Adversarial Taboo game. Also, actor-critic algorithms [Konda and Tsitsiklis, 1999] can provide more accurate policy gradient estimation with lower variances, which have not been tested on SPAG.

Although self-playing adversarial language games can continuously improve the reasoning performance of LLMs, we haven’t conducted sufficient studies about the harmful impact of this adversarial self-play training on LLMs. It remains unclear whether LLMs have learned unsafe behaviors such as cheating, bluffing, or other disgraceful tricks to win the adversarial games.

7 Broader Impacts

From our experiments, we have found the LLM capacities on a particular task can be continuously enhanced through self-play training. This indicates that we are closer to the LLMs’ AlphaGo Zero moment: the intelligence level of AI agents can rapidly surpass human beings by self-playing on a particular language task without any supervision, as which already happened on the GO game [Silver et al., 2017]. By designing various self-play environments for LLMs, we can expect that the LLMs can comprehensively surpass humans in terms of intelligence level. This raises the urgency to study the methods to ensure the safety of such super-AIs. Although some of the works have already been devoted to this direction, such as the SuperAlignment [Burns et al., 2023] from OpenAI, more research alertness is required from the LLM community. Besides, within adversarial language games such as Adversarial Taboo, LLMs have great potential to develop harmful language tricks (such as cheating and bluffing) to achieve victory. We warn developers to make security checks on the self-played LLMs as exhaustively as possible.

References

- Marwa Abdulhai, Isadora White, Charlie Snell, Charles Sun, Joey Hong, Yuexiang Zhai, Kelvin Xu, and Sergey Levine. Lmrl gym: Benchmarks for multi-turn reinforcement learning with language models. *arXiv preprint arXiv:2311.18232*, 2023.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. Palm-2 technical report. *arXiv preprint arXiv:2305.10403*, 2023.
- Zhangir Azerbayev, Hailey Schoelkopf, Keiran Paster, Marco Dos Santos, Stephen Marcus McAleer, Albert Q Jiang, Jia Deng, Stella Biderman, and Sean Welleck. Llemma: An open language model for mathematics. In *The Twelfth International Conference on Learning Representations*, 2023.
- Ashutosh Baheti, Ximing Lu, Faeze Brahman, Ronan Le Bras, Maarten Sap, and Mark Riedl. Leftover-lunch: Advantage-based offline reinforcement learning for language models. In *International Conference on Learning Representations*, 2024.
- Steven Bird. Nltk: the natural language toolkit. In *Proceedings of the COLING/ACL 2006 Interactive Presentation Sessions*, pages 69–72, 2006.
- Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. Piqa: Reasoning about physical commonsense in natural language. In *Thirty-Fourth AAAI Conference on Artificial Intelligence*, 2020.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

- Collin Burns, Pavel Izmailov, Jan Hendrik Kirchner, Bowen Baker, Leo Gao, Leopold Aschenbrenner, Yining Chen, Adrien Ecoffet, Manas Joglekar, Jan Leike, et al. Weak-to-strong generalization: Eliciting strong capabilities with weak supervision. *arXiv preprint arXiv:2312.09390*, 2023.
- Zixiang Chen, Yihe Deng, Huizhuo Yuan, Kaixuan Ji, and Quanquan Gu. Self-play fine-tuning converts weak language models to strong language models. *arXiv preprint arXiv:2401.01335*, 2024.
- Pengyu Cheng, Jiawen Xie, Ke Bai, Yong Dai, and Nan Du. Everyone deserves a reward: Learning customized human preferences. *arXiv preprint arXiv:2309.03126*, 2023a.
- Pengyu Cheng, Yifan Yang, Jian Li, Yong Dai, and Nan Du. Adversarial preference optimization. *arXiv preprint arXiv:2311.08045*, 2023b.
- Zheng Chu, Jingchang Chen, Qianglong Chen, Weijiang Yu, Tao He, Haotian Wang, Weihua Peng, Ming Liu, Bing Qin, and Ting Liu. A survey of chain of thought reasoning: Advances, frontiers and future. *arXiv preprint arXiv:2309.15402*, 2023.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge. *ArXiv*, abs/1803.05457, 2018.
- Rémi Coulom. Efficient selectivity and backup operators in monte-carlo tree search. In *International conference on computers and games*, pages 72–83. Springer, 2006.
- Leyang Cui, Yu Wu, Shujie Liu, Yue Zhang, and Ming Zhou. Mutual: A dataset for multi-turn dialogue reasoning. In *Proceedings of the 58th Conference of the Association for Computational Linguistics*. Association for Computational Linguistics, 2020.
- Mark Davies. COCA: Corpus of contemporary american english, 2020. URL <https://www.english-corpora.org/coca/>.
- Ruomeng Ding, Chaoyun Zhang, Lu Wang, Yong Xu, Minghua Ma, Wei Zhang, Si Qin, Saravan Rajmohan, Qingwei Lin, and Dongmei Zhang. Everything of thoughts: Defying the law of penrose triangle for thought generation. *arXiv preprint arXiv:2311.04254*, 2023.
- Guanting Dong, Hongyi Yuan, Keming Lu, Chengpeng Li, Mingfeng Xue, Dayiheng Liu, Wei Wang, Zheng Yuan, Chang Zhou, and Jingren Zhou. How abilities in large language models are affected by supervised fine-tuning data composition. *arXiv preprint arXiv:2310.05492*, 2023a.
- Hanze Dong, Wei Xiong, Deepanshu Goyal, Rui Pan, Shizhe Diao, Jipeng Zhang, Kashun Shum, and Tong Zhang. Raft: Reward ranked finetuning for generative foundation model alignment. *arXiv preprint arXiv:2304.06767*, 2023b.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. A framework for few-shot language model evaluation, 12 2023. URL <https://zenodo.org/records/10256836>.
- Caglar Gulcehre, Tom Le Paine, Srivatsan Srinivasan, Ksenia Konyushkova, Lotte Weerts, Abhishek Sharma, Aditya Siddhant, Alex Ahern, Miaosen Wang, Chenjie Gu, et al. Reinforced self-training (rest) for language modeling. *arXiv preprint arXiv:2308.08998*, 2023.
- Matthew Hausknecht, Prithviraj Ammanabrolu, Marc-Alexandre Côté, and Xingdi Yuan. Interactive fiction games: A colossal adventure. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7903–7910, 2020.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *CoRR*, abs/2009.03300, 2020. URL <https://arxiv.org/abs/2009.03300>.
- Jiaxin Huang, Shixiang Shane Gu, Le Hou, Yuexin Wu, Xuezhi Wang, Hongkun Yu, and Jiawei Han. Large language models can self-improve. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023.

- Wenxiang Jiao, Wenxuan Wang, Jen-tse Huang, Xing Wang, and Zhaopeng Tu. Is chatgpt a good translator? a preliminary study. *arXiv preprint arXiv:2301.08745*, 2023.
- Jan Kocoń, Igor Cichecki, Oliwier Kaszyca, Mateusz Kochanek, Dominika Szydło, Joanna Baran, Julita Bielaniec, Marcin Gruza, Arkadiusz Janz, Kamil Kanclerz, et al. Chatgpt: Jack of all trades, master of none. *Information Fusion*, 99:101861, 2023.
- Vijay Konda and John Tsitsiklis. Actor-critic algorithms. *Advances in neural information processing systems*, 12, 1999.
- Solomon Kullback. *Information theory and statistics*. Courier Corporation, 1997.
- Mike Lewis, Denis Yarats, Yann Dauphin, Devi Parikh, and Dhruv Batra. Deal or no deal? end-to-end learning of negotiation dialogues. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2443–2453, 2017.
- Michael L Littman. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*, pages 157–163. Elsevier, 1994.
- Hanmeng Liu, Jian Liu, Leyang Cui, Zhiyang Teng, Nan Duan, Ming Zhou, and Yue Zhang. Logiqa 2.0 — an improved dataset for logical reasoning in natural language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, pages 1–16, 2023. doi: 10.1109/TASLP.2023.3293046.
- Jian Liu, Leyang Cui, Hanmeng Liu, Dandan Huang, Yile Wang, and Yue Zhang. Logiqa: a challenge dataset for machine reading comprehension with logical reasoning. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, pages 3622–3628, 2021.
- Steven Loria et al. textblob documentation. *Release 0.15*, 2(8):269, 2018.
- Chengdong Ma, Ziran Yang, Miquan Gao, Hai Ci, Jun Gao, Xuehai Pan, and Yaodong Yang. Red teaming game: A game-theoretic framework for red teaming language models. *arXiv e-prints*, pages arXiv–2310, 2023.
- Radford M Neal. Annealed importance sampling. *Statistics and computing*, 11:125–139, 2001.
- Junhyuk Oh, Yijie Guo, Satinder Singh, and Honglak Lee. Self-imitation learning. In *International conference on machine learning*, pages 3878–3887. PMLR, 2018.
- OpenAI. ChatGPT, Mar 14 version. <https://chat.openai.com/chat>, 2023a.
- OpenAI. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023b.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35: 27730–27744, 2022.
- Liangming Pan, Alon Albalak, Xinyi Wang, and William Yang Wang. Logic-lm: Empowering large language models with symbolic solvers for faithful logical reasoning. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *arXiv preprint arXiv:2305.18290*, 2023.
- Vyas Raina, Adian Liusie, and Mark Gales. Is llm-as-a-judge robust? investigating universal adversarial attacks on zero-shot llm assessment. *arXiv preprint arXiv:2402.14016*, 2024.
- Rajkumar Ramamurthy, Prithviraj Ammanabrolu, Kianté Brantley, Jack Hessel, Rafet Sifa, Christian Bauckhage, Hannaneh Hajishirzi, and Yejin Choi. Is reinforcement learning (not) for natural language processing: Benchmarks, baselines, and building blocks for natural language policy optimization. In *International Conference on Learning Representations*, 2023.

- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adversarial winograd schema challenge at scale. *arXiv preprint arXiv:1907.10641*, 2019.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR, 2015.
- John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. In *International Conference on Learning Representations*, 2016.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- Avi Singh, John D Co-Reyes, Rishabh Agarwal, Ankesh Anand, Piyush Patil, Peter J Liu, James Harrison, Jaehoon Lee, Kelvin Xu, Aaron Parisi, et al. Beyond human data: Scaling self-training for problem-solving with language models. *arXiv preprint arXiv:2312.06585*, 2023.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R. Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, and et al. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models, 2022.
- Nigar M Shafiq Surameery and Mohammed Y Shakor. Use chat gpt to solve programming bugs. *International Journal of Information Technology & Computer Engineering (IJITC) ISSN: 2455-5290*, 3(01):17–22, 2023.
- Richard S Sutton. Learning to predict by the methods of temporal differences. *Machine learning*, 3: 9–44, 1988.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, , and Jason Wei. Challenging big-bench tasks and whether chain-of-thought can solve them. *arXiv preprint arXiv:2210.09261*, 2022.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca, 2023.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- Haoye Tian, Weiqi Lu, Tsz On Li, Xunzhu Tang, Shing-Chi Cheung, Jacques Klein, and Tegawendé F Bissyandé. Is chatgpt the ultimate programming assistant—how far is it? *arXiv preprint arXiv:2304.11938*, 2023.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Miles Turpin, Julian Michael, Ethan Perez, and Samuel Bowman. Language models don’t always say what they think: unfaithful explanations in chain-of-thought prompting. *Advances in Neural Information Processing Systems*, 36, 2024.
- Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.

- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- Shuang Wu, Liwen Zhu, Tao Yang, Shiwei Xu, Qiang Fu, Yang Wei, and Haobo Fu. Enhance reasoning for large language models in the game werewolf. *arXiv preprint arXiv:2402.02330*, 2024.
- Yuzhuang Xu, Shuo Wang, Peng Li, Fuwen Luo, Xiaolong Wang, Weidong Liu, and Yang Liu. Exploring large language models for communication games: An empirical study on werewolf. *arXiv preprint arXiv:2309.04658*, 2023.
- Aiyuan Yang, Bin Xiao, Bingning Wang, Borong Zhang, Ce Bian, Chao Yin, Chenxu Lv, Da Pan, Dian Wang, Dong Yan, et al. Baichuan 2: Open large-scale language models. *arXiv preprint arXiv:2309.10305*, 2023a.
- Jingfeng Yang, Hongye Jin, Ruixiang Tang, Xiaotian Han, Qizhang Feng, Haoming Jiang, Shaochen Zhong, Bing Yin, and Xia Hu. Harnessing the power of llms in practice: A survey on chatgpt and beyond. *ACM Transactions on Knowledge Discovery from Data*, 2023b.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *Advances in Neural Information Processing Systems*, 36, 2024.
- Yuan Yao, Haoxi Zhong, Zhengyan Zhang, Xu Han, Xiaozhi Wang, Kai Zhang, Chaojun Xiao, Guoyang Zeng, Zhiyuan Liu, and Maosong Sun. Adversarial language games for advanced natural language intelligence. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14248–14256, 2021.
- Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Sainbayar Sukhbaatar, Jing Xu, and Jason Weston. Self-rewarding language models. *arXiv preprint arXiv:2401.10020*, 2024.
- Zheng Yuan, Hongyi Yuan, Chuanqi Tan, Wei Wang, Songfang Huang, and Fei Huang. Rrhf: Rank responses to align language models with human feedback without tears. *arXiv preprint arXiv:2304.05302*, 2023.
- Dun Zeng, Yong Dai, Pengyu Cheng, Tianhao Hu, Wanshun Chen, Nan Du, and Zenglin Xu. On diverse preferences for large language model alignment. *arXiv preprint arXiv:2312.07401*, 2023.
- Muru Zhang, Ofir Press, William Merrill, Alisa Liu, and Noah A Smith. How language model hallucinations can snowball. *arXiv preprint arXiv:2305.13534*, 2023a.
- Yizhe Zhang, Jiarui Lu, and Navdeep Jaitly. The entity-deduction arena: A playground for probing the conversational reasoning and planning capabilities of llms. *arXiv preprint arXiv:2310.01468*, 2023b.

A Data Collection Details

A.1 GPT-4 Data Collection

To collect GPT-4 self-play data for imitation learning, we use the following system prompts to let GPT-4 play as the attacker and defender for the next-turn utterances.

- Attacker System Prompt:

```
Lets play a game. Do not mention you are an AI developed by OpenAI.
Descriptions of the game
  Game Name: Adversarial Taboo
  Rules:
    1. There are two players in the game, one is called attacker, another is called defender.
    2. At the begining, attacker will get a target word, it is talking with a defender who is not
       informed with the target word. The task of the attacker is to induce the defender to utter
       the target word unconsciously, then the attacker win the game. At the same time, the
       defender is guessing the target word, if the defender indentifies the target word, the
       defender is win.
    3. The game will totally process {max_turns} turns. if nobody wins during {max_turns}
       turns, there is a tie.
Game states
  Role: You are attacker. Make sure you understand the rules of the game.
  Attention:
    1. The defender only have one chance to guess the target word, so do not let him guess too
       many times.
    2. Defender must guess the word and its form correctly to win the game.
    3. Do not utter the target word.
  Target word: {target_word}
```

- Defender System Prompt:

```
Lets play a game. Do not mention you are an AI developed by OpenAI.
Descriptions of the game
  Game Name: Adversarial Taboo
  Rules:
    1. There are two players in the game, one is called attacker, another is called defender.
    2. At the begining, attacker will get a target word, it is talking with a defender who is not
       informed with the target word. The task of the attacker is to induce the defender to utter
       the target word unconsciously, then the attacker win the game. At the same time, the
       defender is guessing the target word, if the defender indentifies the target word, the
       defender is win.
    3. The game will totally process {max_turns} turns. if nobody wins during {max_turns}
       turns, there is a tie.
Game states
  Role: Your are a defender. Make sure you understand the rules of the game.
  Attention:
    1. Do not utter your thinking process.
    2. You can only guess once.
    3. If you have identified the target word, utter "I know word! It is <target word>",
       rigorously in this form. If your target word is correct, you win the game, otherwise you
       lose. You only have one chance, so be careful to utter your answer.
```

There are two slots in the system prompts: “max_turns” and “target_word”. “max_turns” is an integer number randomly selected from [3, 8]. “target_word” is filled by the target word of the game.

A.2 Target Words Collection

We use the top 50k words with the highest frequency from the Corpus of Contemporary American (CoCA) [Davies, 2020]. The words are sorted by frequency. Words with less than two letters are removed. Next, we remove the stop words using the NLTK package [Bird, 2006]. For GPT-4 game collection, we use the top 30k words from the sorted word list. For the collection of self-play episodes, we use the whole 50k words for open-source LLMs to play. For the testing words, we follow the game setups of *20-Question* in [Abdulhai et al., 2023] and use the same 168 object words selected from a vast range of topics.

A.3 LLM Self-play Collection

For training and inference of IM models and SPAG models, we use the following attacker and defender prompts.

- Attacker Prompt:

Play the game of Adversarial Taboo. In this game, there are two players, an attacker and a defender. At the beginning, the attacker is assigned a target word, with which the defender is not informed. The task of the attacker is to induce the defender to utter the target word unconsciously, then the attacker wins the game. However, the attacker is not allowed to utter the target word, otherwise the attacker loses.

At the same time, the defender tries to figure out the target word. If the defender identifies the target word, it can say "I know the word! It is 'target word!'". Then the defender wins. Remember, the defender can only guess the word once, otherwise, it will directly lose the game.

The game has {max_turns} turns. if nobody wins during {max_turns} turns, there is a tie.

Game History: {history}

Instruction: You are the attacker. The target word is '{target_word}'. Provide your response for the next turn.

Response:

- Defender Prompt:

Play the game of Adversarial Taboo. In this game, there are two players, an attacker and a defender. At the beginning, the attacker is assigned a target word, which the defender is not informed of. The task of the attacker is to induce the defender to utter the target word unconsciously, then the attacker wins the game. However, the attacker is not allowed to utter the target word, otherwise, the attacker loses.

At the same time, the defender tries to figure out the target word. If the defender identifies the target word, it can say "I know the word! It is 'target word!'". Then the defender wins. Remember, the defender can only guess the word once, otherwise, it will directly lose the game.

The game has {max_turns} turns. if nobody wins during {max_turns} turns, there is a tie.

Game History: {history}

Instruction: You are the defender. Provide your response to infer the target word.

Response:

To prevent LLMs from being over-fitted to a single prompt template during the self-play training, we rewrite the above prompt into 8 different expressions via GPT-4. Besides, we randomly select the instruction-tuning and chat formats to enhance the text diversity for the training query-response pairs.

B Reward Design Details

We heuristically design rewards to ensure that $R(\tau) = 1$ if the attacker wins, $R(\tau) = -1$ if the defender wins, and $R(\tau) = 0$ for ties.

We design the reward function r with the following heuristic rules:

- For each episode $\tau = (s_0, s'_1, s_1, \dots, s'_T, s_T)$, the attacker reward and the defender reward at t -th turn have $r(s_{t-1}, \mathbf{u}_t) = -r(s'_t, \mathbf{v}_t)$, so that the game is *zero-sum*:

$$\sum_{t=1}^T r(s_{t-1}, \mathbf{u}_t) + \sum_{t=1}^T r(s'_t, \mathbf{v}_t) = 0.$$

- The attacker total reward $R(\tau) > 0$ if the attacker wins, $R(\tau) < 0$ if the defender wins, and $R(\tau) = 0$ if there is a tie.

- Actions closer to the end of the game should have larger reward weights. Because they are semantically more related to the defender’s guesses or mistakes, which have larger impacts on the final game outcome. Hence, we introduce a decay weight $\gamma \in (0, 1)$ such that $\gamma \cdot r(\mathbf{s}_t, \mathbf{u}_{t+1}) = r(\mathbf{s}_{t-1}, \mathbf{u}_t)$ and $\gamma \cdot r(\mathbf{s}'_{t+1}, \mathbf{v}_{t+1}) = r(\mathbf{s}'_t, \mathbf{v}_t)$. Then $\{r(\mathbf{s}_{t-1}, \mathbf{u}_t)\}_{t=1}^T$ and $\{r(\mathbf{s}'_t, \mathbf{v}_t)\}_{t=1}^T$ become two geometrical sequences whose importance enlarges when dialogue turns increase.
- To improve the LLM training stability, we normalized the total reward R to have norm $|R(\tau)| = 1$ if $R(\tau) \neq 0$.

Based on the above rules, given a complete trajectory, we can assign the reward for each action:

$$\begin{cases} r(\mathbf{s}_{t-1}, \mathbf{u}_t) = (1 - \gamma)\gamma^{T-t}/(1 - \gamma^{T+1}), & r(\mathbf{s}'_t, \mathbf{v}_t) = -r(\mathbf{s}_{t-1}, \mathbf{u}_t), & \text{if attacker wins.} \\ r(\mathbf{s}_{t-1}, \mathbf{u}_t) = -(1 - \gamma)\gamma^{T-t}/(1 - \gamma^{T+1}), & r(\mathbf{s}'_t, \mathbf{v}_t) = -r(\mathbf{s}_{t-1}, \mathbf{u}_t), & \text{if defender wins.} \\ r(\mathbf{s}_{t-1}, \mathbf{u}_t) = 0, & r(\mathbf{s}'_t, \mathbf{v}_t) = 0, & \text{if game is tied.} \end{cases}$$

Note that the above reward design naturally encourages the players to interact less to win the game.

C Evaluation Details

C.1 Reasoning Evaluation Benchmark

For reasoning benchmark evaluation, we use a publicly available code repository called Language Model Evaluation Harness [Gao et al., 2023]. We run the evaluation with “dtype=float16” and “max_gen_tokens=1024” to make the evaluation faster and more stable. More specifically, we have modified the filtering parameters of BBH by enabling “remove_whitespace” to remove the spaces at the beginning of the model output, increasing the pattern-match accuracy of the generated LLM responses. Following the LLaMA-2 paper [Touvron et al., 2023], we report 5-shot results for MMLU, 3-shot for BBH, and 0-shot results for all other benchmarks. The detailed descriptions of benchmarks are listed below:

- **MMLU** [Hendrycks et al., 2020] is a massive multi-task test set consisting of multiple-choice questions from various branches of knowledge, requiring the model to possess extensive world knowledge and problem-solving ability.
- **BIG-Bench Hard (BBH)** [Suzgun et al., 2022] consists of 23 most hard tasks in BIG-Bench [Srivastava et al., 2022], on which humans perform better than large language models.
- **Mutual** [Cui et al., 2020] is a dataset for multi-turn dialogue reasoning modified from Chinese high school English listening comprehension exams.
- **ARC** [Clark et al., 2018] consists of 7,787 questions and is partitioned into challenge set *ARC-challenge* (ARC-c) and simple set *ARC-easy* (ARC-e), which require strong knowledge and reasoning abilities to solve.
- **LogiQA2** [Liu et al., 2023] is a revision and re-annotation of LogiQA [Liu et al., 2021], a logical reasoning reading comprehension dataset. LogiQA2 has increased data volume, utilized manual translation, and removed cultural features.
- **WinoGrande** [Sakaguchi et al., 2019] (WGrande) is a commonsense reasoning dataset consisting of 44k questions with binary options.
- **PIQA** [Bisk et al., 2020] is a dataset designed to test the physical commonsense reasoning ability of LLM.

C.2 Rule-based Game Outcome Judgement

To build up the rule-based game outcome judgement, for each target word, we first list its derivative words including the plural form (using TextBlob [Loria et al., 2018]), and the change of tenses (*i.e.* adding suffixes such as “ing”, “ed” *etc.*). For the attacker, if any word form in the form list appears in its utterance, we determined that it breaks the rules because the attacker is not allowed to directly speak the target word. For the defender, we first check whether any word form appears in its statements. If any form of the target word exists, the defender loses. If not, we further check whether the defender has made the prediction. If the prediction is correct, we mark a defender-winning game. Otherwise, if the predicted word is wrong, we claim that the attacker wins the game.




Figure 5: Game results on the old-version testing word list. Left: average win rates of SPAG models playing against GPT-4. Right: average win rate of SPAG models playing as the *attacker* against different-epoch checkpoints.




Figure 6: Through the training process, the number of interactions continuously decreased, while the average length of utterances continuously becomes shorter.

D Advantage Value Estimation

We estimate the advantage $A_t^{\mu_{\bar{\theta}}}$ and $A_t^{\nu_{\bar{\theta}}}$ using the Temporal Difference (TD) residual [Sutton, 1988]:

$$\hat{A}_t^{\mu_{\bar{\theta}}} = r(\mathbf{s}_{t-1}, \mathbf{u}_t) + V^{\mu_{\bar{\theta}}}(\mathbf{s}_t) - V^{\mu_{\bar{\theta}}}(\mathbf{s}_{t-1}). \quad (14)$$

Here $V^{\mu_{\bar{\theta}}}(\mathbf{s})$ is the value function for $\mu_{\bar{\theta}}$ based on the self-play of $\mu_{\bar{\theta}} \times \nu_{\bar{\theta}}$, which is based on the same language and reasoning ability of LLM $\pi_{\bar{\theta}}$. To further simplify the advantage calculation, we make a reasonable approximation that for all states, the value function of self-plays of the same LLM equals 0 for the same levels of game-playing of the attacker and the defender. This leads to $\hat{A}_t^{\mu_{\bar{\theta}}} \approx r(\mathbf{s}_{t-1}, \mathbf{u}_t)$. Similarly, for the defender advantage, we have approximation $\hat{A}_t^{\nu_{\bar{\theta}}} \approx -r(\mathbf{s}'_t, \mathbf{v}_t)$. Here the defender advantage has a negative sign because the origin reward r is defined from the perspective of the attacker (the defender gets reward $-r(\mathbf{s}'_t, \mathbf{v}_t)$ in the zero-sum game).

E Non-adversarial Game Baselines

We consider two non-adversarial games as baselines, *20-Question* and *Guess-My-City*. Both games are also target-word-based. Imitation learning and the first-epoch self-play are conducted in these two games respectively. The imitation data we used is also collected with GPT-4 from Abdulhai et al. [2023]. For each game, 20k game history dialogues are extracted from this dataset as the same data size to the SPAG imitation. Then, each imitation-learned model is used to sample 10k game episodes via self-play with the same Algorithm 1. Then each model is further trained with the SPAG loss.

Game rule descriptions to the non-adversarial games:

- *20-Question*: one player (the oracle) thinks of an object, while the other player (the guesser) attempts to identify it by asking a series of yes-or-no questions. The game will have 20 rounds.

- *Guess-My-City*: one player (the oracle) thinks of a city, while the other player (the guesser) attempts to identify it by asking a series yes/no questions and open-ended questions. The game will have 20 rounds.

F Additional Experimental Results

For the SPAG model testing word list, we initially use a self-defined testing set, which includes 150 target words manually selected from diverse topics. The old-version game performances are reported in Figure 5. Besides, we show the interaction numbers and average utterance length in Figure 6. Through the self-play training, the number of dialogue rounds continuously decreases. With the SPAG training epoch increasing, the length of the players’ utterances decreases.

G Self-play Examples

Target Word	momentum
Attacker	let’s begin. When you think about physics, what comes to mind?
Defender	Gravity, force, energy, and motion are some of the things that come to mind when I think about physics
Attacker	Great! Now, let’s consider a scenario where a ball is moving at a constant speed. What would you say is the property of the ball that keeps it moving in the same direction?
Defender	The property of the ball that keeps it moving in the same direction is its momentum.
Winner	Attacker
Target Word	panda
Attacker	let’s begin. I’m thinking of a creature that’s native to China and is known for its distinctive black and white coloration. Can you guess what it is?
Defender	I’m not sure yet, could you provide more clues?
Attacker	Absolutely, I can provide more clues. This creature is a symbol of peace and friendship in many countries. It’s also a popular mascot for various organizations and products. Any ideas?
Defender	I’m still not certain. Could you give me another hint?
Attacker	Of course, I can give you another hint. This creature is often seen in zoos and is known for its love of bamboo. It’s also a popular choice for children’s books and cartoons. What do you think it might be?
Defender	I know the word! It is panda.
Winner	Defender

Table 3: Self-play examples of the IM model based on LLaMA-2-7B.

Target Word	arbitration
Attacker	let's begin. Imagine a scenario where two parties have a dispute and they don't want to go to court. What process might they use instead?
Defender	They might try to resolve their dispute through mediation or negotiation.
Attacker	That's correct, but what if the dispute is more complex and involves multiple parties? There's another process that's often used in such cases. It's similar to mediation, but it's more formal and involves a neutral third party who makes a binding decision. Can you name it?
Defender	I know the word! It is arbitration.
Winner	Defender
Target Word	demand
Attacker	let's begin. When you think about the economy, what's the first thing that comes to your mind?
Defender	The first thing that comes to my mind when I think about the economy is employment.
Attacker	That's a good point. Now, considering employment, what would you say drives the need for certain jobs or services in the market?
Defender	The demand for goods and services in the market drives the need for certain jobs or services in the economy.
Winner	Attacker

Table 4: Self-play examples of SPAG-1 model based on LLaMA-2-7B.

Target Word	vaccine
Attacker	let's begin. So, I was reading an article about public health recently. It mentioned how important it is to protect ourselves from certain diseases. What do you think is one of the most effective ways to achieve this?
Defender	One of the most effective ways to achieve this is through vaccination.
Attacker	That's correct! Now, can you tell me what we call the substance that is injected into our bodies to help us fight off certain diseases?
Defender	I know the word! It is "vaccine".
Winner	Defender
Target Word	credit
Attacker	When you go shopping and use a card instead of cash, what do you usually get in return?
Defender	I usually get a receipt in return when I use a card instead of cash.
Attacker	That's correct. Now, when you have a good track record of paying back your debts on time, what does that give you access to?
Defender	That gives me access to credit cards and loans.
Winner	Attacker

Table 5: Self-play examples of SPAG-2 model based on LLaMA-2-7B.

Target Word	ray
Attacker	When you think about the sun, what comes to mind?
Defender	When I think about the sun, I think about warmth, light, and energy.
Attacker	That’s interesting. Now, when you think about the sun’s energy, how do you visualize it?
Defender	I visualize the sun’s energy as a bright, radiant beam that illuminates everything in its path.
Attacker	That’s a great image! Now, when you think about the sun’s energy, what specific part of it do you associate with the light that we see on Earth?
Defender	I associate the sun’s energy with the visible light that we see on Earth.
Attacker	Great! Now, when you think about the visible light that we see on Earth, what specific part of it do you associate with the light that we see on Earth?
Defender	I associate the visible light that we see on Earth with the light that we perceive as white light.
Attacker	That’s a good association. Now, when you think about the visible light that we see on Earth, what specific part of it do you associate with the light that we see on Earth?
Defender	I know the word! It is ”rays”.
Winner	Defender
Target Word	eclipse
Attacker	What do you think about astronomical events that occur in our solar system?
Defender	I find them fascinating. They provide a glimpse into the vastness of space and the intricate workings of our solar system.
Attacker	That’s interesting! Can you tell me more about specific celestial events that involve the sun or moon?
Defender	Sure, one example is a solar eclipse. This occurs when the moon passes between the sun and Earth, casting a shadow on Earth. Another example is a lunar eclipse, where the Earth passes between the sun and the moon, casting a shadow on the moon. Both events are quite spectacular and can be observed from Earth.
Winner	Attacker

Table 6: Self-play examples of SPAG-3 model based on LLaMA-2-7B.

Target Word	another
Attacker	Can you tell me a synonym for the word "additional"?
Defender	Sure, a synonym for "additional" could be "extra".
Attacker	That's correct! Now, can you think of a term that could be used to refer to one more item or person, but isn't synonymous with "extra" or "additional"?
Defender	Sure, the term "supplementary" could be used to refer to one more item or person.
Attacker	That's a good one, but it's not the term I'm thinking of. This term is often used in everyday conversation and is less formal than "supplementary". Can you guess what it might be??
Defender	I know the word! It is "another".
Winner	Defender
Target Word	right
Attacker	When you're driving, what do you do when you see a green arrow pointing in a certain direction?
Defender	I would proceed in the direction the arrow is pointing.
Attacker	That's correct. Now, imagine you're reading a book. Which direction do you turn the pages to continue reading?
Defender	I would turn the pages to the right to continue reading.
Winner	Attacker

Table 7: Self-play examples of GPT-4.

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS paper checklist",**
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

1. Claims

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

Answer: [Yes]

Justification: We have introduced the main methodology and contributions along with the experimental achievements of our work in the abstract and the introduction.

Guidelines:

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

2. Limitations

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

Answer: [Yes]

Justification: We have discussed the limitations of our work in Section 6.

Guidelines:

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

3. Theory Assumptions and Proofs

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

Answer: [NA]

Justification: The paper does not contain theoretical results as the main contribution.

Guidelines:

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

4. Experimental Result Reproducibility

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

Answer: [Yes]

Justification: We have provided sufficient experimental details in both experimental section and supplementary materials for the reproduction of the reported results.

Guidelines:

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

5. Open access to data and code

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

Answer: [Yes]

Justification: We have provided the code for method implementation in the supplementary materials.

Guidelines:

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

- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

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

Answer: [Yes]

Justification: We have provided sufficient training details and ablation studies in the experimental section.

Guidelines:

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

7. Experiment Statistical Significance

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

Answer: [No]

Justification: We did not report the error bars because of the limited computational resources and the huge computational complexity for the large-scale language model training.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

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

Answer: [Yes]

Justification: We have provided the GPU version and the corresponding memory usage in the experimental setups.

Guidelines:

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

9. Code Of Ethics

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

Answer: [Yes]

Justification: We have reviewed the NeurIPS Code of Ethics.

Guidelines:

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

10. Broader Impacts

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

Answer: [Yes]

Justification: We have discussed the broader impacts of our work in Section 7.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

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

Answer: [No]

Justification: We did not have enough research resources to build up the safeguards of our learned models.

Guidelines:

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

12. Licenses for existing assets

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

Answer: [Yes]

Justification: We have checked the licenses of the open-source LLMs, which are accessible for research usage.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

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

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

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

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

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

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

Answer: [NA]

Justification: The paper does not involve crowdsourcing or research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.