EvoBench: Towards Real-world LLM-Generated Text Detection Benchmarking for Evolving Large Language Models

Anonymous ACL submission

Abstract

With the widespread of Large Language Models (LLMs), there has been an increasing need to detect LLM-generated texts, prompting extensive research in this area. However, existing detection methods mainly evaluate on static benchmarks, which neglect the evolving nature of LLMs. Relying on existing static benchmarks could create a misleading sense of security, overestimating the real-world effectiveness of detection methods. To bridge this gap, we introduce EvoBench, a dynamic benchmark considering a new dimension of generalization across continuously evolving LLMs. EvoBench categorizes the evolving LLMs into (1) updates over time and (2) developments like finetuning and pruning, covering 7 LLM families and their 30 evolving versions. To measure the generalization across evolving LLMs, we introduce a new EMG (Evolving Model Generalization) metric. Our evaluation of 14 detection methods on EvoBench reveals that they all struggle to maintain generalization when confronted with evolving LLMs. To mitigate the generalization problems, we further propose improvement strategies, demonstrating EMG performance improvements up to 12.15%. Our research sheds light on critical challenges in real-world LLM-generated text detection and represents a significant step toward practical applications.[1]

1 Introduction

Large Language Models (LLMs), such as ChatGPT (OpenAI, 2022b), Claude (Anthropic, 2024), and LLaMA (Touvron et al., 2023a), have demonstrated remarkable capabilities in natural language understanding and task processing, leading to their widespread application (M Alshater, 2022; Yuan et al., 2022; Christian, 2023). However, concerns have been raised about the misuse of these models in areas like social media (Ahmed et al., 2021), education (Lee et al., 2023), and academic writing (Mitchell, 2022; Patrick Wood, 2023). For instance, LLMs can be used to manipulate the public by generating comments or to fabricate experimental data and statistical results in support of unverified hypotheses (Solaiman et al., 2019; Goldstein et al., 2023). The potential misuse of LLMs highlights the urgent need to detect LLM-generated text (Kaur et al., 2022; Chen and Shu, 2023).

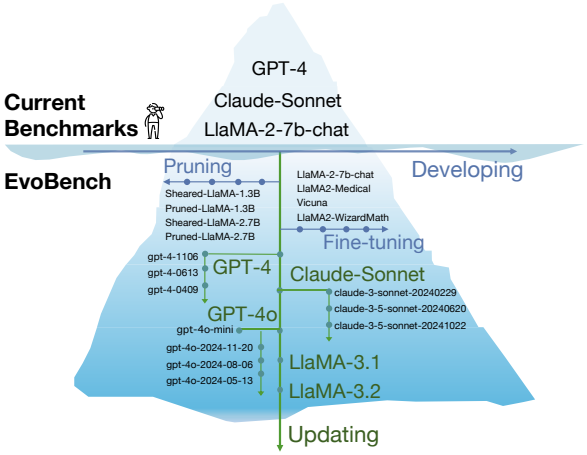

Figure 1: Current benchmarks (Guo et al., 2023; Bao et al.; Wang et al., 2024b,a; Kwan et al., 2024) primarily focus on specific versions of LLMs and neglect vast ecosystem of evolving LLMs beneath the surface, including updates over time or developments through fine-tuning or pruning. The figure primarily includes LLMs before Jan 2025.

The academic community has carried out extensive research to detect LLM-generated text effectively (Liu et al., 2019; Gehrmann et al., 2019; Su et al., 2023; Solaiman et al., 2019). Current methods can be categorized into supervised methods (Hu et al., 2023; Yu et al., 2024; Chen et al., 2024b) and zero-shot methods (Ippolito et al., 2020; Yang et al.; Mitchell et al., 2023; Su et al., 2023; Bao et al.). Supervised methods are typically trained with a binary classifier to distinguish be-

[1]The Evobench is now available at: https://anonymous.4open.science/r/EvoBench.

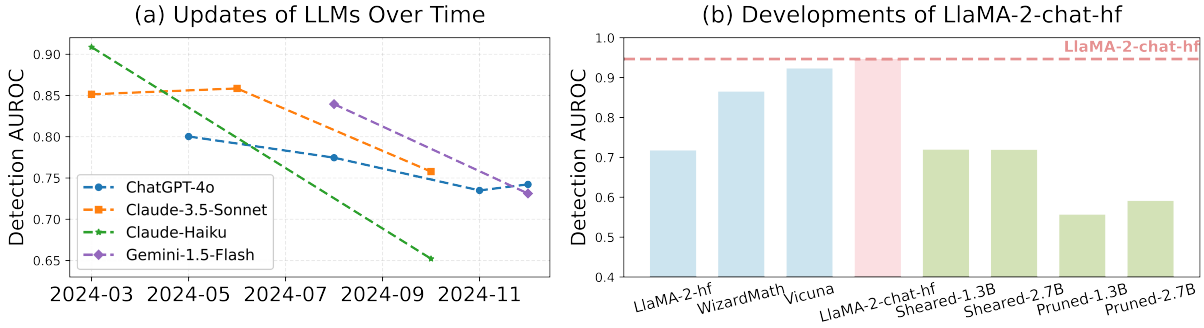

Figure 2: Detection accuracy (measured in AUROC) of Fast-DetectGPT when faced with the evolving LLMs. Figure (a) shows a clear decline in average detection performance as the LLM updates. Figure (b) illustrates the developments of LLMs, including fine-tuning and pruning.

tween texts generated by LLMs and those created by humans, while zero-shot methods primarily rely on statistical features gathered from pre-trained large language models. Although these methods show strong performance on existing benchmarks (Bao et al.; Hu et al., 2023; Chen et al., 2024b; Yu et al., 2024), they often fall short in real-world applications as current benchmarks neglect the evolving nature of LLMs (Wang et al., 2024b; Guo et al., 2023; Wang et al., 2024a; He et al., 2024; Macko et al., 2023). As illustrated in Figure 1, existing benchmarks include a limited versions of LLMs, much like observers seeing only the tip of an iceberg while overlooking the vast evolving LLMs hidden beneath the surface.

In real-world applications, LLMs are continuously evolving via regular updates, fine tuning (Touvron et al., 2023b; Zhang et al., 2023), or pruning (Sun et al.; Liang et al., 2021), all of which affect LLMs' output (Tao et al., 2024; Touvron et al., 2023a; Gunasekar et al., 2023), thereby impacting the detection performance. Figure 2 illustrates that the detection performance of current detection methods suffers up to 25% drop as LLMs evolve.[2] Therefore, relying on current static benchmarks for evaluation could create a misleading sense of security, leading the research community to overestimate the real-world effectiveness of detection methods. Therefore, it is crucial to enhance existing benchmarks to better capture the ongoing evolving LLMs.

In this paper, we propose EvoBench, a dynamic benchmark that extends traditional benchmarking to account for the evolving LLMs. EvoBench aims to provide a more accurate evaluation of detection methods by incorporating two dimensions of evolving LLMs: **updates** and **developments**, as shown in Figure 1. Updates refer to changes made by the LLM publishers, including LLM updates, *i.e.*, GPT-4o undergoes updates approximately every 3 months [3]. On the other hand, developments refer to optimizations made by LLM developers for specific application scenarios, such as fine-tuning and pruning. EvoBench introduces a new dimension of generalization, enabling the evaluation of detection methods across evolving LLMs, covering 7 widely used LLMs and their 30 evolving versions.

Using EvoBench, we evaluate 14 widely used detection methods and find that all struggle to adapt to the evolving nature of LLMs. To intuitive quantify the generalization across evolving LLMs, we introduce the EMG (Evolving Model Generalization) metric. Results reveal that the performance of 14 current detection methods, including widely used Fast-DetectGPT (Bao et al.) and RADAR (Hu et al., 2023) detectors, significantly drop when faced with evolving versions, as reflected by low EMG values, while some achieve high AUROC (0.91) scores on individual models[4]. This phenomenon highlights the limitations of current detection methods in generalization capabilities.

To mitigate this problem, we further explore two strategies. For zero-shot methods like Fast-DetectGPT (Bao et al.), we propose to prune the scoring model to extract shared features for detection, which could lead to a 12.15% improvement in EMG performance across developments of the LLM. For supervised methods, improving the quality and distribution of training data - by contin-

[2]The widely used detection method, Fast-DetectGPT, declines up to 25% drop when detecting the Claude-3-5-haiku-20241022 compared to Claude-3-haiku-20240307.

[3]These versions are gpt-4o-2024-05-13, gpt-4o-2024-08-06, gpt-4o-2024-11-20, and chatgpt-4o-latest.

[4]Fast-DetectGPT achieve 0.91 AUROC when detecting Claude-Haiku-2024-03-07, but drop to 0.65 AUROC when detecting the version of 2024-10-22.

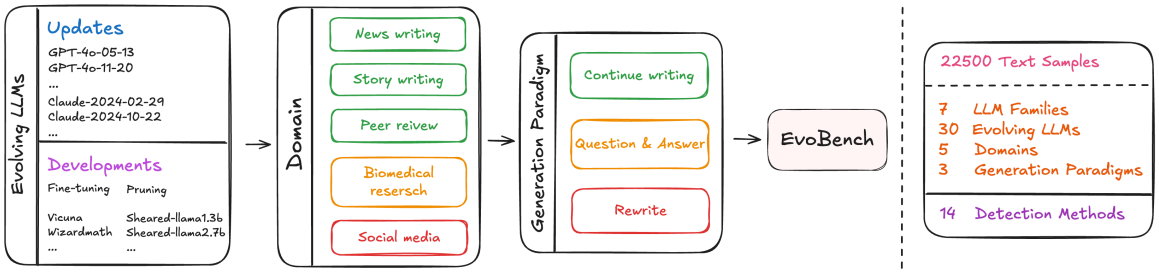

Figure 3: The construction pipeline and the dataset statistics of EvoBench. EvoBench includes three dimensions of generalization: dataset domains, generation strategies, and evolving LLMs. The evolving LLMs are further divided into two categories: updates and developments.

uously incorporating data from evolving LLMs - could increase 0.24% EMG performance across updates of the GPT-4. However, ensuring effective generalization across both updates and developments simultaneously remains a significant challenge. This finding provides valuable insight into the complexities of adapting detection methods to the continuously evolving LLMs.

Overall, EvoBench serves as a complementary benchmark that focuses on the evolving nature of LLMs, offering a more accurate framework for evaluating the real-world applicability of detection methods. While it does not claim to be a one-size-fits-all solution, we envision EvoBench evolving alongside detection methods and other benchmarks, helping guide future developments in this area. In particular, we suggest that extracting shared features from evolving LLMs could enhance the generalization of detection methods. With EvoBench, the community can more effectively evaluate and refine detection methods, contributing to more robust and practical detection systems that mitigate LLM misuse in real-world applications.

2 EvoBench

2.1 Definition of Evolving LLMs

EvoBench primarily considers two evolving pathways for LLMs: (1) **updates**, which are released by publishers, and (2) **developments**, which involve optimization made by developers. These two evolving pathways present significant challenges to current detection methods.

Updates. We focus on two types of updates that occur over time: model (Brown et al., 2020; Tao et al., 2024) and version updates (Brown et al., 2020). Model updates typically involve significant changes to the model's architecture. For example, the transition from GPT-4 to GPT-4o represents a major architectural shift (Achiam et al., 2023). In contrast, version updates occur more frequently, with shorter cycles and less perception, such as the regular releases of new versions of GPT-4 every 2-3 months, such as updates in May, August, and November 2024.

Developments. In EvoBench, model developments include fine-tuning (Hu et al., 2022; Mangrulkar et al., 2022), pruning (Liang et al., 2021), often used for domain adaptation and optimizing efficiency for real-world applications. Unlike updates, these actions are nearly imperceptible to the detector (Bhattacharjee et al., 2023). For example, developers may fine-tune LLMs using private datasets, making the detection more challenging.

2.2 Dataset Collection

2.2.1 Data Collection Principle

EvoBench, incorporating the two evolving pathways of LLMs, extends existing benchmarks (Mitchell et al., 2023) widely used in current detection methods, including DetectGPT (Mitchell et al., 2023), Fast-DetectGPT (Bao et al.), and ImBD (Chen et al., 2024b) to assess three key dimensions of generalization: (1) evolving LLMs, (2) domain, and (3) generation paradigms. Figure 3 illustrates our data collection pipeline to cover these three key dimensions.

Generalization to LLM Evolution. To comprehensively assess evolving, we consider 7 widely used LLM families: GPT-4o, GPT-4, Claude, Gemini, Qwen, LlaMA3, and LlaMA2, encompassing a total of 30 evolving versions. For updates, we focus on families of GPT-4o, GPT-4, Claude, Gemini, Qwen, and LlaMA3, tracking how detection methods perform across updates of these models. For developments, we focus on the LlaMA2 family, incorporating fine-tuning and pruning techniques. Detailed specifications of all LLMs are provided in

the Appendix C.

Generalization to Domains. EvoBench includes five datasets spanning diverse tasks: *XSum* (Narayan et al., 2018) for (1)news articles, *WritingPrompts* (Fan et al., 2018) for (2) creative story writing, and *PubMed* (Jin et al., 2019) for (3) biomedical research question answering. Additionally, we also include two datasets: *SocialMedia* (Kula and Gregor, 2024) for (4) datasets on social media and *PeerRead* (Kang et al., 2018) for (5) peer-reviewed academic writing.

Generalization to Generation Paradigms. EvoBench incorporates three distinct generation paradigms: (1) continuation-based generation, includes *XSum*, *Writing*, and *PeerRead*; (2) question-answering generation, including *PubMed*; (3) paraphrased generation, including *SocialMedia*. By integrating these diverse paradigms, EvoBench offers a more comprehensive framework to assess the robustness of detection methods in various text generation paradigms.

2.2.2 Dataset Collection Process

We detail our systematic dataset construction process following the three-dimensional generalization framework. For each dataset, we first carefully selected 150 human-authored texts as reference samples. To collect AI-generated texts, we employed consistent prompting strategies across all evaluated models. Specifically, for continuation-based generation tasks (*Xsum*, *Writing*, *PeerRead*), we provided the prefix of human-written texts as context. For *PubMed*, we maintained the original question-answer format, while for *SocialMedia*, we implemented a paraphrasing approach where models were tasked with preserving semantic meaning while using different words. Details of prompts are shown in Appendix E.

2.3 Dataset Statistics

EvoBench consists of 7 LLM families, encompassing a total of 30 evolving versions. For each evolving LLM, we collect samples across 5 distinct datasets, with each dataset containing a balanced distribution of 150 human-authored texts and 150 machine-generated texts. In total, EvoBench comprises $30 \times 5 \times 150 = 22500$ machine-generated samples, resulting in a comprehensive evaluation benchmark of $22,500$ pairs of text samples.

2.4 Evaluation Metrics

To quantify the generalization ability of detection methods to evolving LLMs, we introduce the EMG Φ_E (Evolving Model Generalization) metric, inspired by the coefficient of variation. It evaluates the consistency and trend of detection performance across evolving LLMs.

For an LLM family with m evolving versions, the widely evaluated version is selected as the $base$ model, and the remaining $m-1$ versions as evolving models. The performance (measured by AUROC Φ_A) change $\Delta\Phi_A^i$ between the i-th models and the $base$ model is :

$$\Delta\Phi_A^i = \Phi_A^i - \Phi_A^{base}, \tag{1}$$

where Φ_A^i and Φ_A^{base} is the Φ_A of the i-th evolving model and $base$ model, respectively.

The average performance change μ_E of $\Delta\Phi_A$, representing the overall performance change, is:

$$\mu_E = \frac{1}{m-1}\sum_{i=1}^{m-1}\left(\Delta\Phi_A^i\right). \tag{2}$$

The volatility σ_E of $\Delta\Phi_A$ using the standard deviation is:

$$\sigma_E = \sqrt{\frac{1}{m-1}\sum_{i=1}^{m-1}\left(\Delta\Phi_A^i - \mu_E\right)^2}. \tag{3}$$

The proposed Φ_E is defined as follows:

$$\Phi_E = \frac{\mu_E}{\sigma_E + \lambda} \times \gamma, \tag{4}$$

where λ is a regularization term to prevent fluctuations, and γ is a scaling factor. In this study, we set λ and γ as 1 and 100, respectively.

2.5 Detection Methods

We evaluate a range of supervised and zero-shot detection methods to assess their generalization ability against evolving LLMs. Additionally, we conducted repeated generation processes to ensure that output diversity does not significantly impact detection performance in Appendix B.

2.5.1 Existing Methods

For supervised detectors, we tested the GPT-2 detectors developed by OpenAI (Liu et al., 2019), the RADAR detector (Hu et al., 2023), Text Fluoroscopy (Yu et al., 2024), and ImBD (Imitate Before Detect) (Chen et al., 2024b). For zero-shot detectors, we included Likelihood (average

log probabilities) (Gehrmann et al., 2019), LRR (a hybrid method combining log probability and log-rank) (Su et al., 2023), LogRank (mean of the log ranks sorted in descending order of probabilities) (Solaiman et al., 2019), Entropy (average token-level entropy from the predictive distribution)(Ippolito et al., 2020), DNA-GPT (Yang et al.), DetectGPT (Mitchell et al., 2023), and its enhanced variants NPR (Su et al., 2023) and Fast-DetectGPT (Bao et al.).

2.5.2 Optimizing Strategies

Enhancing Supervised Detectors via optimizing Training Data. To enhance detection generalization across evolving models, we propose an incremental training approach: For a evolving LLM, we generate additional training data D_{evolving} and update the training set D_{old}:

$$D_{\text{new}} = D_{\text{old}} \cup D_{\text{evolving}}. \tag{5}$$

Using the new dataset D_{new}, the detector f_{detector} is retrained to adapt to the evolving LLM.

Enhancing Zero-shot Detector via Extracting Shared Features. To enhance the detection generalization of the zero-shot method, we propose to prune the scoring model p_θ^{pruned} to extract the shared features from the developments of LLMs. Based on Fast-DetectGPT, we define our detection metric $d(x, p_\theta^{\text{pruned}})$ based on conditional probability curvature:

$$d(x, p_\theta^{\text{pruned}}) = \frac{\log p_\theta^{\text{pruned}}(x|x) - \tilde{\mu}}{\tilde{\sigma}}, \tag{6}$$

where $\tilde{\mu}$ represents the expected log probability under the pruned model:

$$\tilde{\mu} = \mathbb{E}_{\tilde{x}\sim p_\theta^{\text{pruned}}(\tilde{x}|x)}[\log p_\theta^{\text{pruned}}(\tilde{x}|x)], \tag{7}$$

and $\tilde{\sigma}^2$ captures the variance of these log probabilities:

$$\tilde{\sigma}^2 = \mathbb{E}_{\tilde{x}\sim p_\theta^{\text{pruned}}(\tilde{x}|x)}\left[\left(\log p_\theta^{\text{pruned}}(\tilde{x}|x) - \tilde{\mu}\right)^2\right]. \tag{8}$$

3 Exprimental Results

3.1 Setup

Detection Methods. To mimic real-world scenarios, we set up a black-box environment in which the detector is assumed to be unaware of the source model of the text to be detected under evaluation. We tested 14 detection methods, including 5 supervised detectors and 9 zero-shot detectors. Among them, RADAR (Hu et al., 2023) and Fast-DetectGPT (Bao et al.) are well-known methods in their respective categories. Details of experimental settings can be found in the Appendix D and Appendix F.

3.2 Main Results

To begin with, we select two leading detection methods, Fast-DetectGPT and RADAR detector, and evaluate them on our EvoBench, the results are shown in Table 1. The evaluation results of other detection methods are shown in Appendix H.

First, a trend is that texts generated by **more advanced LLMs are harder to detect in the same LLM family.** For instance, comparing Claude-3-Haiku[5] and Claude-3-Opus [6], the latter demonstrates more advanced capabilities. When detecting these two LLMs, Fast-DetectGPT and RADAR showed decreased detection AUROC of 3.58% and 2.77%, respectively, as shown in Table 1. Similarly, comparing detecting *SocialMedia* datasets generated by GPT-4o-mini with GPT-4o, two detectors showed a decline in detection AUROC of 3.22% and 6.69%, respectively.

Our results also reveal an intriguing pattern: **existing detection methods often over-optimize for specific LLM versions**, rather than maintaining robust generalization capabilities across different evolving LLMs. Specifically, detectors with high AUROC on widely used LLM families often show performance degradation with newer versions of the same family. For example, Fast-DetectGPT performed excellently on the GPT-4o-05-13 version, achieving an average detection accuracy of 0.8003 on EvoBench, but dropped to 0.7422 when facing the GPT-4o-latest version. Similar patterns were observed in other widely used LLM families, such as GPT-4 and Claude-Sonnet. In contrast, when evaluating less widely used LLMs like LlaMA3, detectors tend to demonstrate better generalization capabilities across evolving LLMs. For example, the detection performance of the RADAR remained stable across different versions of the LlaMA3 family, with AUROC scores across five datasets ranging from 0.7989 to 0.8348.

This suggests that the essence of **generalization lies in the overfitting problem**. Current detectors are often optimized to achieve superior performance on existing benchmarks, which leads to

[5] claude-3-haiku-20240307
[6] claude-3-opus-20240229

LLMs	Version Time/ Version Name	Fast-DetectGPT						RADAR					
		Xsum	Writing	PubMed	SocialMedia	PeerRead	Avg.	Xsum	Writing	PubMed	SocialMedia	PeerRead	Avg.
Updates													
GPT-4o	2024-05-13	0.90	0.97	0.73	0.58	0.83	0.80	0.99	0.83	0.83	0.60	0.68	0.79
	2024-08-06	0.87 (−2.75%)	0.94 (−2.41%)	0.70 (−3.54%)	0.59 (+1.76%)	0.77 (−5.89%)	0.77 (−2.57%)	0.99 (−0.11%)	0.86 (+3.27%)	0.82 (−0.53%)	0.62 (+2.12%)	0.69 (+0.75%)	0.80 (+1.10%)
	2024-11-20	0.72 (−17.6%)	0.90 (−6.90%)	0.70 (−3.77%)	0.57 (−0.09%)	0.78 (−4.25%)	0.73 (−6.54%)	0.99 (−0.87%)	0.72 (−11.0%)	0.79 (−4.08%)	0.62 (+2.02%)	0.66 (−1.47%)	0.75 (−3.09%)
	Latest	0.70 (−20.0%)	0.91 (−6.13%)	0.70 (−3.36%)	0.58 (+0.64%)	0.82 (−0.14%)	0.74 (−5.81%)	0.99 (−0.79%)	0.75 (−7.49%)	0.80 (−2.83%)	0.64 (+4.46%)	0.67 (−0.53%)	0.77 (−1.44%)
GPT-4o-mini	2024-07-18	0.91 (+0.91%)	0.97 (+0.12%)	0.73 (−0.42%)	0.61 (+3.02%)	0.80 (−2.14%)	0.80 (+0.30%)	1.00 (+0.33%)	0.87 (+3.82%)	0.82 (−0.68%)	0.66 (+6.69%)	0.64 (−3.96%)	0.80 (+1.24%)
GPT-4	2023-06-13	0.93	0.98	0.76	0.49	0.75	0.78	0.99	0.87	0.89	0.66	0.75	0.83
	2023-11-06	0.92 (−0.85%)	0.93 (−4.76%)	0.73 (−2.94%)	0.55 (+5.43%)	0.83 (+7.71%)	0.79 (+0.92%)	0.99 (−0.01%)	0.79 (−8.49%)	0.80 (−8.81%)	0.66 (+0.24%)	0.63 (−12.3%)	0.77 (−5.88%)
	2024-01-25	0.89 (−3.22%)	0.94 (−3.41%)	0.70 (−5.75%)	0.54 (+4.23%)	0.82 (+7.34%)	0.78 (−0.16%)	0.99 (−0.23%)	0.80 (−7.23%)	0.84 (−5.09%)	0.65 (−0.46%)	0.66 (−8.68%)	0.79 (−4.34%)
	2024-04-09	0.83 (−9.30%)	0.91 (−6.56%)	0.72 (−4.39%)	0.58 (+8.27%)	0.79 (+4.08%)	0.77 (−1.58%)	0.99 (−0.36%)	0.79 (−8.45%)	0.82 (−7.07%)	0.60 (−5.92%)	0.60 (−14.8%)	0.76 (−7.32%)
Claude-Sonnet	2024-02-29	0.95	0.98	0.81	0.58	0.93	0.85	0.97	0.86	0.84	0.74	0.61	0.80
	2024-06-20	0.97 (+2.33%)	0.99 (+0.80%)	0.77 (−4.32%)	0.66 (+7.69%)	0.90 (−2.90%)	0.86 (+0.72%)	0.99 (+1.61%)	0.81 (−5.21%)	0.82 (−2.44%)	0.72 (−1.05%)	0.55 (−5.57%)	0.78 (−2.53%)
	2024-10-22	0.90 (−4.43%)	0.93 (−5.39%)	0.68 (−13.9%)	0.51 (−7.80%)	0.78 (−15.3%)	0.76 (−9.37%)	0.95 (−1.52%)	0.72 (−13.4%)	0.74 (−9.85%)	0.67 (−6.74%)	0.55 (−6.38%)	0.73 (−7.59%)
Claude-Haiku	2024-03-07	1.00	1.00	0.86	0.75	0.94	0.91	1.00	0.93	0.84	0.77	0.67	0.84
	2024-10-22	0.84 (−15.5%)	0.92 (−7.56%)	0.64 (−21.7%)	0.39 (−35.9%)	0.46 (−47.4%)	0.65 (−25.6%)	1.00 (−0.07%)	0.86 (−7.50%)	0.73 (−10.9%)	0.70 (−6.38%)	0.82 (+14.7%)	0.82 (−2.04%)
Claude-Opus	2024-02-29	0.97 (−2.49%)	0.96 (−3.89%)	0.82 (−3.80%)	0.72 (−3.27%)	0.89 (−4.46%)	0.87 (−3.58%)	0.99 (−0.33%)	0.87 (−6.01%)	0.82 (−2.23%)	0.77 (+0.29%)	0.62 (−5.57%)	0.81 (−2.77%)
Qwen	Qwen1.5-7B	0.92	0.99	0.66	0.49	0.91	0.80	0.93	0.90	0.72	0.51	0.48	0.71
	Qwen2-7B	0.99 (+6.86%)	1.00 (+0.29%)	0.74 (+7.35%)	0.46 (−2.75%)	0.89 (−1.37%)	0.82 (+2.08%)	0.89 (−4.13%)	0.87 (−3.42%)	0.82 (+9.46%)	0.52 (+1.35%)	0.54 (+6.14%)	0.73 (+1.88%)
	Qwen2.5-7B	0.99 (+6.51%)	1.00 (+0.10%)	0.79 (+12.3%)	0.47 (−2.11%)	0.90 (−0.78%)	0.83 (+3.22%)	0.90 (−2.60%)	0.91 (+0.66%)	0.79 (+6.85%)	0.52 (+1.31%)	0.53 (+4.95%)	0.73 (+2.23%)
LlaMA3	LlaMA-3.1-8B	0.99	0.98	0.85	0.69	0.97	0.90	1.00	0.89	0.78	0.72	0.63	0.80
	LlaMA-3.1-70B	1.00 (+0.31%)	1.00 (+1.50%)	0.81 (−3.72%)	0.67 (−1.73%)	0.96 (−0.96%)	0.89 (−0.92%)	1.00 (−0.19%)	0.88 (−1.32%)	0.76 (−1.57%)	0.74 (+1.29%)	0.62 (−1.22%)	0.80 (−0.60%)
	LlaMA-3.2-1B	0.96 (−3.46%)	0.99 (+1.09%)	0.92 (+7.29%)	0.86 (+17.5%)	0.96 (−0.30%)	0.94 (+4.43%)	0.96 (−3.51%)	0.94 (+4.39%)	0.82 (+3.74%)	0.75 (+2.80%)	0.71 (+7.52%)	0.83 (+2.99%)
	LlaMA-3.2-3B	0.99 (−0.13%)	0.99 (+0.31%)	0.91 (+6.56%)	0.77 (+8.32%)	0.96 (−0.30%)	0.93 (+2.95%)	0.99 (−0.44%)	0.92 (+2.71%)	0.82 (+3.70%)	0.77 (+4.23%)	0.66 (+2.80%)	0.83 (+2.60%)
	LlaMA-3.3-70B	1.00 (+0.32%)	1.00 (+1.61%)	0.79 (−5.53%)	0.69 (+0.15%)	0.90 (−6.38%)	0.88 (−1.97%)	1.00 (−0.32%)	0.89 (−0.74%)	0.79 (+0.94%)	0.73 (+0.76%)	0.65 (+2.29%)	0.81 (+0.59%)
Developments													
Fine-tuning	LlaMA-2-7B-chat-hf	0.97	0.99	0.91	0.90	0.97	0.95	0.68	0.63	0.49	0.64	0.64	0.62
	Vicuna-1.5-7B	1.00 (+2.36%)	1.00 (+0.52%)	0.80 (−10.4%)	0.88 (−1.58%)	0.94 (−2.54%)	0.92 (−2.34%)	0.97 (+29.3%)	0.93 (+29.7%)	0.81 (+31.5%)	0.65 (+0.94%)	0.64 (−0.25%)	0.80 (+18.2%)
	Wizardmath-7B	0.92 (−5.17%)	0.98 (−1.00%)	0.74 (−17.0%)	0.85 (−4.73%)	0.84 (−12.8%)	0.87 (−8.14%)	0.75 (+6.68%)	0.76 (+12.5%)	0.70 (+20.5%)	0.69 (+4.61%)	0.54 (−9.33%)	0.69 (+7.01%)
Pruning	Sheared-LlaMA1.3B	0.65 (−31.7%)	0.89 (−10.1%)	0.48 (−42.0%)	0.86 (−3.29%)	0.71 (−26.2%)	0.72 (−22.7%)	0.77 (+8.98%)	0.66 (+2.96%)	0.59 (+9.45%)	0.73 (+8.59%)	0.54 (−10.2%)	0.66 (+3.96%)
	Sheared-LlaMA1.3B-pruned	0.41 (−56.0%)	0.81 (−18.2%)	0.35 (−55.3%)	0.75 (−14.3%)	0.46 (−51.0%)	0.56 (−38.9%)	0.92 (+24.0%)	0.79 (+15.7%)	0.70 (+20.8%)	0.87 (+22.7%)	0.65 (+1.59%)	0.79 (+16.9%)
	Sheared-LlaMA2.7B-pruned	0.45 (−52.4%)	0.87 (−12.3%)	0.37 (−53.9%)	0.75 (−14.2%)	0.52 (−44.9%)	0.59 (−35.5%)	0.84 (+15.7%)	0.76 (+12.8%)	0.59 (+9.59%)	0.78 (+14.3%)	0.55 (−9.25%)	0.70 (+8.66%)
	Sheared-LlaMA2.7B	0.70 (−26.7%)	0.92 (−7.50%)	0.50 (−40.0%)	0.82 (−7.72%)	0.65 (−31.7%)	0.72 (−22.7%)	0.75 (+6.88%)	0.68 (+4.81%)	0.58 (+8.83%)	0.67 (+3.47%)	0.47 (−17.2%)	0.63 (+1.35%)

Table 1: The detection performance (measured in AUROC) of two leading detection methods, Fast-DetectGPT and RADAR, on EvoBench. For each LLM family, we set the LLMs that are widely evaluated in text generation detection benchmarks as anchor points. 'Latest' refers to the LLM currently used in the web version of GPT-4o.

overfitting to a specific LLM version rather than being truly effective. This also reveals the vulnerability of current detection methods across evolving LLMs. From a practical perspective, a key priority for future research lies in developing more adaptable detection frameworks that can simultaneously preserve detection performance while demonstrating robust generalization across dataset domains and evolving LLMs.

3.2.1 Detection Generalization across Evolving LLMs

To visually demonstrate the generalization ability of detection methods across evolving LLMs, we evaluate the 14 detection methods using the EMG metric, with results shown in Table 2.

Table 2 demonstrates that no method can maintain stable performance across all evolving LLMs. Specifically, detectors like ImBD and Text-Fluoroscopy perform well on complex models like GPT-4o and GPT-4. Meanwhile, detectors like entropy and RADAR exhibit better generalizations of newer model families, including Qwen, LlaMA3, and the developments of LlaMA2. This variation in performance across different detectors suggests that improving generalization might require adopting strategies that combine the strengths of different approaches, which may help improve the adaptability of the detector and reduce sensitivity to specific LLM evolution trends.

3.2.2 Performance of Optimizing Strategies

To improve the generalization of current detection methods across evolving LLMs, we have pre-

Methods	Updates						Developments		Avg.
	GPT-4o	GPT-4	Claude-Sonnet	Claude-haiku	Qwen	LlaMA3	Fine-tuning	Pruning	
Supervised Detectors									
RoBERTa-base	−2.252	−4.593	−8.325	−15.95	+5.862	+1.457	−2.378	+6.871	−2.413
RoBERTa-large	−1.878	−5.577	−10.06	−14.63	+4.646	+3.448	−0.725	+12.54	−1.530
RADAR	−0.537	−5.775	−4.934	−2.394	+2.053	+1.372	+11.96	+7.302	+1.131
Text-Fluoroscopy	−0.493	−0.655	+0.902	−0.123	−5.241	+1.895	−11.03	−15.77	−3.815
Imitate Before Detect	+0.440	+2.397	−1.572	−4.445	+0.023	−0.054	+8.942	−9.842	−0.513
Zero-Shot Detectors									
Likelihood	−0.873	−0.047	−3.810	−8.280	−1.174	+0.439	−6.991	−32.99	−6.717
Rank	−2.127	+1.071	−0.178	−12.09	+4.249	+1.023	−9.661	−7.109	−3.103
LogRank	−1.316	−0.106	−4.272	−9.443	−0.608	+0.837	−6.851	−30.14	−6.488
Entropy	−3.202	+0.416	+0.602	−3.411	+4.138	+2.145	+5.150	+25.12	+3.869
LRR	−2.511	−0.075	−5.864	−14.26	+2.629	+1.911	−6.964	−17.45	−5.325
NPR	−2.846	+7.436	−2.237	−15.26	+7.132	+0.866	−8.063	−31.18	−5.520
DNA-GPT	−3.805	+1.249	−4.621	−12.26	−1.332	+1.337	−6.355	−36.11	−7.738
DetectGPT	−4.095	+10.04	+2.258	−11.09	+2.357	−0.336	−5.442	−20.18	−3.312
Fast-DetectGPT	−3.558	−0.271	−4.115	−13.16	+2.632	+1.094	−5.095	−27.94	−6.302

Table 2: EMG performance of 14 detection methods on EvoBench. Red indicates a negative EMG, signifying a decrease in AUROC when facing evolving LLMs, while a larger EMG value reflects a better generalization of the detection method.

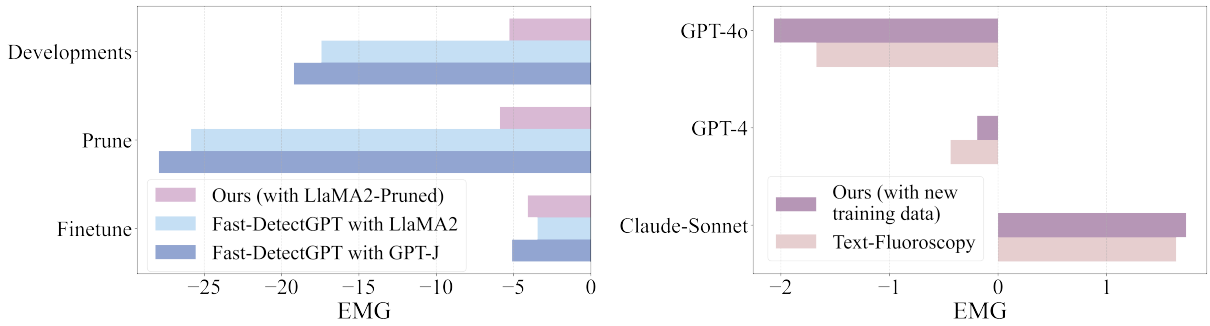

Figure 4: EMG performance of two optimizing strategies compared with their corresponding original methods. The left panel illustrates the improvement for zero-shot detection using the pruned scoring model, while the right panel shows the optimization for supervised detection using the newer dataset.

liminarily explored two possible approaches, with results presented in Figure 4. For the zero-shot detection method, we compared the EMG performance of the Fast-DetectGPT (Bao et al.) detector with different scoring models, including GPT-J, LlaMA2, and the pruned version of LlaMA2. The results show that our proposed method outperforms by 12.152% in EMG when dealing with developments, including fine-tuning and pruning. However, for fine-tuned developments, using the pruned LlaMA2 sacrifices some performance compared to directly using LlaMA2.

For the supervised detection method, we chose to improve Text-Fluoroscopy (Yu et al., 2024). The original training dataset was generated by GPT-3.5-turbo. We enhanced it by regenerating it with the initial version of the corresponding LLM family and retrained the detector. The details are shown in the Appendix D. We observed improvements in performance for GPT-4 and Claude-Sonnet but a decline for GPT-4o. This suggests that while dataset enhancement can improve performance in some LLMs, it may not always yield better results when faced with different evolving LLMs.

4 Related Work

4.1 Evolving LLMs

Updates of Pre-trained LLMs. The advent of LLMs, such as ChatGPT (OpenAI, 2022a) and GPT-4 (OpenAI, 2023), has marked a paradigm shift in text generation. However, these pre-trained models are not static (Tao et al., 2024; Touvron et al., 2023a; Gunasekar et al., 2023; Biderman et al., 2023); rather, they undergo continuous evo-

lution, with models like ChatGPT frequently updating their parameters to enhance performance and adapt to new application scenarios (Zheng et al., 2024; Touvron et al., 2023a).

Developments on Pre-trained LLMs. The rise of open-source LLMs has facilitated their widespread usage in diverse domains. Developers can flexibly tailor LLMs to specific tasks through fine-tuning (Touvron et al., 2023a,b; Zhang et al., 2023), pruning (Sun et al.; Liang et al., 2021; Xia et al.), and quantization (Du et al., 2024; Chen et al., 2024a; Dettmers et al., 2023) techniques that modify the model's parameters or structure, ultimately influencing output characteristics. In this process, developers could use domain-specific datasets to enhance the model's understanding of specialized fields and their contextual knowledge (Kerner, 2024). Furthermore, depending on practical application scenarios, users may select models with varying parameter scales to strike an optimal balance between performance, speed, and resource consumption (Kim et al., 2024; Zhao et al., 2023).

4.2 LLM-generated Text Detection

Supervised Methods. Supervised methods (Liu et al., 2019; Chen et al., 2024b) are usually trained to differentiate between texts generated by LLMs and texts created by humans. For example, RADAR (Hu et al., 2023) introduces the idea of adversarial learning to train a detector that can resist paraphrase attacks. Text Fluoroscopy (Yu et al., 2024) extracts discriminative features from the intermediate layers of the language model and utilizes them to train a binary classifier, which enhances the generalization of supervised detectors across texts from different semantic domains. However, as LLMs continue to evolve, frequent adaptation mechanisms—including updates, fine-tuning, pruning, and other optimization strategies—introduce changes in their generated text, making it difficult to guarantee the effectiveness of existing detection methods on newer model versions. This creates a significant gap between academic detection models and real-world applications, ultimately limiting the practical utility of these methods.

Zero-shot Methods. Existing zero-shot methods primarily rely on statistical features extracted using pre-trained large language models (Bao et al.; Mitchell et al., 2023). These features include likelihood (Gehrmann et al., 2019), probability curvature (Mitchell et al., 2023), divergence between multiple completions of a truncated passage (Yang et al.), and conditional probability curvature (Bao et al.). Zero-shot detection methods are immune to domain-specific degradation, demonstrating superior generalization in detection tasks (Gehrmann et al., 2019; Mitchell et al., 2023). However, the effectiveness of zero-shot detection heavily depends on the alignment between the pre-trained LLM used for detection and the generation model that produced the text to be detected (Bao et al.; Mitchell et al., 2023). While current methods achieve state-of-the-art results on existing benchmarks (Guo et al., 2023; Bao et al.; Wang et al., 2024b,a; Kwan et al., 2024), the evolving nature of LLMs introduces significant shifts in this alignment. As models continue to evolve, the pre-trained LLMs used for detection may become increasingly misaligned with the updated generation models, leading to a decline in detection performance. This explains why many methods, despite ranking highly on leaderboards, often fail when deployed in real-world scenarios.

5 Conclusion

In this paper, we introduce EvoBench, a novel benchmark for evaluating the generalization of LLM-generated text detection methods across evolving LLMs. EvoBench defines two key dimensions of LLM evolution: (1) updates made by LLM publishers over time and (2) developments carried out by developers, ensuring a comprehensive understanding of how evolving LLMs impact detection performance. To quantify this evolving dimension generalization, we propose the EMG (Evolving Model Generalization) metric.

We evaluate 14 widely used detection methods using EvoBench, revealing their vulnerabilities when facing evolving LLMs. In response, to improve the generalization of zero-shot methods across developing LLMs, we propose two strategies: for zero-shot methods, we suggest pruning the scoring model to extract shared features across LLM developments; for supervised methods, we recommend augmenting training data with data from LLM updates. Our benchmark represents a key step towards bringing detection methods into real-world scenarios for evolving LLMs. We also envision continuously updating this benchmark to cover a broader range of evolving LLMs, enabling more comprehensive evaluations across various domains and evolving LLMs.

Limitations

Although we have thoroughly considered the three key dimensions of generalization toward real-world scenarios, certain limitations remain that warrant further exploration in future research. First, it is important to emphasize that EvoBench is not intended to replace existing benchmarks for evaluating the generalization of LLM-generated text detection methods but rather to complement them. Therefore, EvoBench primarily focuses on three key dimensions in real-world scenarios: domains, generation strategies, and evolving models. We have not covered other aspects of generalization, such as language. A potential direction for future research is to extend EvoBench to incorporate additional dimensions of generalization.

Second, robustness has not been addressed in this study, as our focus was mainly on detection generalization. However, EvoBench could serve as a foundation for inspiring current robustness research, as robustness also involves tasks such as rewriting text using different LLMs.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Alim Al Ayub Ahmed, Ayman Aljabouh, Praveen Kumar Donepudi, and Myung Suh Choi. 2021. Detecting fake news using machine learning: A systematic literature review. *arXiv preprint arXiv:2102.04458*.

AI Anthropic. 2024. Claude 3.5 sonnet model card addendum. *Claude-3.5 Model Card*, 3(6).

Guangsheng Bao, Yanbin Zhao, Zhiyang Teng, Linyi Yang, and Yue Zhang. Fast-detectgpt: Efficient zero-shot detection of machine-generated text via conditional probability curvature. In *The Twelfth International Conference on Learning Representations*.

Amrita Bhattacharjee, Tharindu Kumarage, Raha Moraffah, and Huan Liu. 2023. ConDA: Contrastive domain adaptation for AI-generated text detection. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 598–610, Nusa Dua, Bali. Association for Computational Linguistics.

Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O'Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, et al. 2023. Pythia: A suite for analyzing large language models across training and scaling. In *International Conference on Machine Learning*, pages 2397–2430. PMLR.

Sid Black, Gao Leo, Phil Wang, Connor Leahy, and Stella Biderman. 2021. GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Canyu Chen and Kai Shu. 2023. Combating misinformation in the age of llms: Opportunities and challenges. *arXiv preprint arXiv: 2311.05656*.

Hong Chen, Chengtao Lv, Liang Ding, Haotong Qin, Xiabin Zhou, Yifu Ding, Xuebo Liu, Min Zhang, Jinyang Guo, Xianglong Liu, and Dacheng Tao. 2024a. DB-LLM: Accurate dual-binarization for efficient LLMs. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 8719–8730, Bangkok, Thailand. Association for Computational Linguistics.

Jiaqi Chen, Xiaoye Zhu, Tianyang Liu, Ying Chen, Xinhui Chen, Yiwen Yuan, Chak Tou Leong, Zuchao Li, Tang Long, Lei Zhang, et al. 2024b. Imitate before detect: Aligning machine stylistic preference for machine-revised text detection. *arXiv preprint arXiv:2412.10432*.

Jon Christian. 2023. Cnet secretly used ai on articles that didn't disclose that fact, staff say. *Futurusm, January*.

Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. Qlora: Efficient finetuning of quantized llms. *Advances in neural information processing systems*, 36:10088–10115.

DaYou Du, Yijia Zhang, Shijie Cao, Jiaqi Guo, Ting Cao, Xiaowen Chu, and Ningyi Xu. 2024. BitDistiller: Unleashing the potential of sub-4-bit LLMs via self-distillation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 102–116, Bangkok, Thailand. Association for Computational Linguistics.

Angela Fan, Mike Lewis, and Yann Dauphin. 2018. Hierarchical neural story generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898.

Sebastian Gehrmann, Hendrik Strobelt, and Alexander M Rush. 2019. Gltr: Statistical detection and visualization of generated text. In *Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics (ACL).

Josh A Goldstein, Girish Sastry, Micah Musser, Renee DiResta, Matthew Gentzel, and Katerina Sedova. 2023. Generative language models and automated influence operations: Emerging threats and potential mitigations. *arXiv preprint arXiv:2301.04246*.

Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, et al. 2023. Textbooks are all you need. *arXiv preprint arXiv:2306.11644*.

Biyang Guo, Xin Zhang, Ziyuan Wang, Minqi Jiang, Jinran Nie, Yuxuan Ding, Jianwei Yue, and Yupeng Wu. 2023. How close is chatgpt to human experts? comparison corpus, evaluation, and detection. *arXiv preprint arXiv:2301.07597*.

Xinlei He, Xinyue Shen, Zeyuan Chen, Michael Backes, and Yang Zhang. 2024. Mgtbench: Benchmarking machine-generated text detection. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, pages 2251–2265.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.

Xiaomeng Hu, Pin-Yu Chen, and Tsung-Yi Ho. 2023. Radar: Robust ai-text detection via adversarial learning. *Advances in Neural Information Processing Systems*, 36:15077–15095.

Daphne Ippolito, Daniel Duckworth, Chris Callison-Burch, and Douglas Eck. 2020. Automatic detection of generated text is easiest when humans are fooled. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1808–1822.

Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. Pubmedqa: A dataset for biomedical research question answering. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2567–2577.

Dongyeop Kang, Waleed Ammar, Bhavana Dalvi, Madeleine van Zuylen, Sebastian Kohlmeier, Eduard Hovy, and Roy Schwartz. 2018. A dataset of peer reviews (PeerRead): Collection, insights and NLP applications. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1647–1661, New Orleans, Louisiana. Association for Computational Linguistics.

Davinder Kaur, Suleyman Uslu, Kaley J Rittichier, and Arjan Durresi. 2022. Trustworthy artificial intelligence: a review. *ACM Computing Surveys (CSUR)*, 55(2):1–38.

Tobias Kerner. 2024. Domain-specific pretraining of language models: A comparative study in the medical field. *arXiv preprint arXiv:2407.14076*.

Sanghoon Kim, Dahyun Kim, Chanjun Park, Wonsung Lee, Wonho Song, Yunsu Kim, Hyeonwoo Kim, Yungi Kim, Hyeonju Lee, Jihoo Kim, Changbae Ahn, Seonghoon Yang, Sukyung Lee, Hyunbyung Park, Gyoungjin Gim, Mikyoung Cha, Hwalsuk Lee, and Sunghun Kim. 2024. SOLAR 10.7B: Scaling large language models with simple yet effective depth up-scaling. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 6: Industry Track)*, pages 23–35, Mexico City, Mexico. Association for Computational Linguistics.

Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Sebastian Kula and Michal Gregor. 2024. Multilingual models for check-worthy social media posts detection. *arXiv preprint arXiv:2408.06737*.

Wai-Chung Kwan, Xingshan Zeng, Yufei Wang, Yusen Sun, Liangyou Li, Yuxin Jiang, Lifeng Shang, Qun Liu, and Kam-Fai Wong. 2024. M4LE: A multi-ability multi-range multi-task multi-domain long-context evaluation benchmark for large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15568–15592, Bangkok, Thailand. Association for Computational Linguistics.

Jooyoung Lee, Thai Le, Jinghui Chen, and Dongwon Lee. 2023. Do language models plagiarize? In *Proceedings of the ACM Web Conference 2023*, pages 3637–3647.

Tailin Liang, John Glossner, Lei Wang, Shaobo Shi, and Xiaotong Zhang. 2021. Pruning and quantization for deep neural network acceleration: A survey. *Neurocomputing*, 461:370–403.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Muneer M Alshater. 2022. Exploring the role of artificial intelligence in enhancing academic performance: A case study of chatgpt. *Available at SSRN*.

Dominik Macko, Robert Moro, Adaku Uchendu, Jason Lucas, Michiharu Yamashita, Matúš Pikuliak, Ivan Srba, Thai Le, Dongwon Lee, Jakub Simko, et al. 2023. Multitude: Large-scale multilingual machine-generated text detection benchmark. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9960–9987.

Sourab Mangrulkar, Sylvain Gugger, Lysandre Debut, Younes Belkada, Sayak Paul, and Benjamin

Bossan. 2022. Peft: State-of-the-art parameter-efficient fine-tuning methods. https://github.com/huggingface/peft.

Alex Mitchell. 2022. Professor catches student cheating with chatgptl: 'i feel abject terror'. Retrieved from https://nypost.com/2022/12/26/students-using-chatgpt-to-cheat-professor\-warns/. Accessed on February 17, 2023.

Eric Mitchell, Yoonho Lee, Alexander Khazatsky, Christopher D Manning, and Chelsea Finn. 2023. Detectgpt: Zero-shot machine-generated text detection using probability curvature. In *International Conference on Machine Learning*, pages 24950–24962. PMLR.

Shashi Narayan, Shay B Cohen, and Mirella Lapata. 2018. Don't give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807.

OpenAI. 2022a. ChatGPT. https://chat.openai.com/.

OpenAI. 2022b. Introducing chatgpt.

OpenAI. 2023. GPT-4 Technical Report. *arXiv preprint arXiv:2303.08774*.

Mary Louise Kelly Patrick Wood. 2023. 'everybody is cheating': Why this teacher has adopted an open chatgpt policy. Retrieved from https://www.npr.org/2023/01/26/1151499213/chatgpt-ai-education-cheating-classroom\-wharton-school. Accessed on January 26, 2023.

Irene Solaiman, Miles Brundage, Jack Clark, Amanda Askell, Ariel Herbert-Voss, Jeff Wu, Alec Radford, Gretchen Krueger, Jong Wook Kim, Sarah Kreps, et al. 2019. Release strategies and the social impacts of language models. *arXiv preprint arXiv:1908.09203*.

Jinyan Su, Terry Zhuo, Di Wang, and Preslav Nakov. 2023. Detectllm: Leveraging log rank information for zero-shot detection of machine-generated text. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 12395–12412.

Mingjie Sun, Zhuang Liu, Anna Bair, and J Zico Kolter. A simple and effective pruning approach for large language models. In *The Twelfth International Conference on Learning Representations*.

Zhengwei Tao, Ting-En Lin, Xiancai Chen, Hangyu Li, Yuchuan Wu, Yongbin Li, Zhi Jin, Fei Huang, Dacheng Tao, and Jingren Zhou. 2024. A survey on self-evolution of large language models. *arXiv preprint arXiv:2404.14387*.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Ben Wang and Aran Komatsuzaki. 2021. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. https://github.com/kingoflolz/mesh-transformer-jax.

Yuxia Wang, Jonibek Mansurov, Petar Ivanov, Jinyan Su, Artem Shelmanov, Akim Tsvigun, Osama Mohammed Afzal, Tarek Mahmoud, Giovanni Puccetti, Thomas Arnold, Alham Aji, Nizar Habash, Iryna Gurevych, and Preslav Nakov. 2024a. M4GT-bench: Evaluation benchmark for black-box machine-generated text detection. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3964–3992, Bangkok, Thailand. Association for Computational Linguistics.

Yuxia Wang, Jonibek Mansurov, Petar Ivanov, Jinyan Su, Artem Shelmanov, Akim Tsvigun, Chenxi Whitehouse, Osama Mohammed Afzal, Tarek Mahmoud, Toru Sasaki, et al. 2024b. M4: Multi-generator, multi-domain, and multi-lingual black-box machine-generated text detection. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1369–1407.

Mengzhou Xia, Tianyu Gao, Zhiyuan Zeng, and Danqi Chen. Sheared llama: Accelerating language model pre-training via structured pruning. In *The Twelfth International Conference on Learning Representations*.

Xianjun Yang, Wei Cheng, Yue Wu, Linda Ruth Petzold, William Yang Wang, and Haifeng Chen. Dna-gpt: Divergent n-gram analysis for training-free detection of gpt-generated text. In *The Twelfth International Conference on Learning Representations*.

Xiao Yu, Kejiang Chen, Qi Yang, Weiming Zhang, and Nenghai Yu. 2024. Text fluoroscopy: Detecting llm-generated text through intrinsic features. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 15838–15846.

Ann Yuan, Andy Coenen, Emily Reif, and Daphne Ippolito. 2022. Wordcraft: story writing with large language models. In *27th International Conference on Intelligent User Interfaces*, pages 841–852.

Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, et al. 2023. Instruction tuning for large language models: A survey. *arXiv preprint arXiv:2308.10792*.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.

Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyan Luo, Zhangchi Feng, and Yongqiang Ma. 2024. Llamafactory: Unified efficient fine-tuning of 100+ language models. *arXiv preprint arXiv:2403.13372*.

A Examples of Texts Generated by Different LLMs

In this section, we provided examples of responses from different updates of the same LLMs to the same question, offering a visual representation of how changes in versions lead to variations in the model's output in Table 3 and Table 4.

B Additional Experiment

Additionally, to eliminate the influence of output diversity on detection performance, we first generated 150 texts from each model for every dataset, resulting in a total of 300 texts. Additionally, we conducted repeated experiments with several models to ensure that output diversity does not significantly impact changes in detection difficulty. By doing so, we ensure that the variations in detection difficulty primarily stem from model evolution itself, rather than biases introduced by the diversity of the generated texts. The results of this experimental process are presented in Table 5, where we highlight the potential range of biases caused by the diversity of generated texts. This helps further to understand the relationship between model evolution and detection performance.

C Details of Evolving LLMs

The details of evolving LLMs are shown in Table 6.

D Details of Optimizing Strategies

In this section, we provide the details of two optimizing strategies. For the zero-shot detection method, we chose to optimize the Fast-DetectGPT detector. We use the pruned model 'princeton-nlp/Sheared-LLaMA-2.7B-Pruned' to replace the scoring model 'meta-llama/Llama-2-7b-hf' model as the scoring model. For the supervised detector, we chose to optimize the Text Fluoroscopy. Following Text Fluoroscopy, we used the first 200 entries of the open-source Human-ChatGPT Comparison Corpus (HC3) (Guo et al., 2023) dataset collected by previous researchers as a training set. The ratio for splitting the training and validation is $8:1$. We use a specific LLM of the LLM families to regenerate machine text in the training dataset and then use the new dataset and the original dataset to retrain the detector. Specifically, we regenerated this part of the training set using GPT-4-0613 and tested the results on the GPT-4 family. For the GPT-4o

Table 3: An example of content generated by different versions of the GPT-4o and GPT-4o-mini

Does airway surgery lower serum lipid levels in obstructive sleep apnea patients?	
Human Answer	Surgery may improve blood lipid levels in patients with OSA, especially in patients with preoperative dyslipidemia, potentially yielding a major benefit in metabolism and cardiovascular sequelae. Prospective studies should examine this ...
GPT-4o 2024-05-13	Airway surgery for obstructive sleep apnea (OSA) has been shown to reduce serum lipid levels in some patients. The improvement is likely due to better oxygenation and reduced inflammation following successful surgery, ...
GPT-4o 2024-08-06	Airway surgery can potentially lower serum lipid levels in obstructive sleep apnea (OSA) patients. Surgical interventions that alleviate OSA symptoms may improve metabolic profiles, including lipid levels, due to reduced hypoxia and ...
GPT-4o 2024-11-20	Yes, airway surgery in obstructive sleep apnea (OSA) patients may lower serum lipid levels. Studies suggest that by improving oxygenation, reducing inflammation, and addressing the metabolic impacts of OSA, surgeries like ...
GPT-4o Latest	Yes, airway surgery for obstructive sleep apnea (OSA) can lower serum lipid levels. By improving airway patency and reducing hypoxia, surgery may enhance lipid metabolism and decrease systemic inflammation, which are linked ...
GPT-4o-mini 2024-07-18	Airway surgery in obstructive sleep apnea patients may lead to improvements in overall health and potentially lower serum lipid levels. However, the extent of lipid level changes can vary among individuals and ...

Table 4: An example of content generated by different versions of the Claude-Haiku and Calude-Opus

Neck pain treatment with acupuncture: does the number of needles matter?	
Human Answer	For this pathology, the number of needles, 5 or 11, seems not to be an important variable in determining the therapeutic effect when the time of stimulation is the same in the two groups.
Claude-Haiku 2024-03-07	Acupuncture for neck pain can be effective, but the number of needles used may not be the most critical factor. The placement and selection of acupuncture points tailored to the individual's condition are more important ...
Claude-Haiku 2024-10-22	Research suggests that the number of acupuncture needles may not significantly impact neck pain treatment. Effectiveness depends more on precise needle placement, practitioner skill, and individual patient response. Further studies are needed to definitively determine ...
Claude-Opus 2024-02-29	The number of needles used in acupuncture for neck pain treatment may impact its effectiveness. However, factors such as needle placement, depth of insertion, and stimulation technique are likely more important than the specific number ...

family, we used GPT-4o-2024-05-13, and for the Claude-Haiku family, we used Claude-3-sonnet-20240229. we use `gte-Qwen1.5-7B-instruct`[7] as the encoder and the classifier consists of three fully connected layers with Tanh function. The dimensions of the intermediate layers in the classifier are 1024 and 512, respectively. The batch size is set to 16, and Adam (Kingma and Ba, 2014) optimizer is employed with an initial learning rate of $3e-3$.

[7]https://huggingface.co/Alibaba-NLP/gte-Qwen1.5-7B-instruct

E Data Collection.

In this section, we describe the process of data collection. Most of these datasets were generated

	Fast-DetectGPT	RADAR	Text-Fluoroscopy	Imitate Before Detect	RoBERTa-base	RoBERTa-large	Likelihood	Rank	LogRank	Entropy	DetectGPT	LRR	NPR	DNA-GPT
GPT-3.5-turbo-2024-01-25	0.8301	0.8280	0.8904	0.8992	0.6621	0.6723	0.7701	0.6432	0.7663	0.4163	0.4910	0.7098	0.6072	0.8160
	0.8420	0.8319	0.8917	0.9011	0.6629	0.6696	0.7822	0.6511	0.7787	0.4118	0.5008	0.7120	0.6276	0.8266
	(+1.19%)	(+0.39%)	(+0.13%)	(+0.19%)	(+0.08%)	(−0.26%)	(+1.21%)	(+0.78%)	(+1.24%)	(−0.44%)	(+0.98%)	(+0.22%)	(+2.05%)	(+1.06%)
GPT-4o-mini-2024-07-18	0.8053	0.7954	0.8800	0.8831	0.5539	0.5799	0.7238	0.6474	0.7133	0.4617	0.5374	0.6341	0.6045	0.7255
	0.8033	0.7976	0.8777	0.8886	0.5566	0.5827	0.7264	0.6494	0.7154	0.4543	0.5417	0.6356	0.6101	0.7291
	(−0.20%)	(+0.22%)	(−0.23%)	(+0.55%)	(+0.27%)	(+0.28%)	(+0.26%)	(+0.20%)	(+0.21%)	(−0.74%)	(+0.43%)	(+0.15%)	(+0.56%)	(+0.36%)
GPT-4-turbo-2024-04-09	0.7644	0.7459	0.8535	0.8534	0.5053	0.5146	0.6959	0.6315	0.6856	0.4526	0.5229	0.6182	0.5982	0.7114
	0.7661	0.7595	0.8504	0.8559	0.5034	0.5099	0.6953	0.6278	0.6848	0.4533	0.5208	0.6109	0.5925	0.6952
	(+0.17%)	(+1.36%)	(−0.31%)	(+0.25%)	(−0.18%)	(−0.47%)	(−0.06%)	(−0.37%)	(−0.08%)	(+0.07%)	(−0.20%)	(−0.72%)	(−0.57%)	(−1.63%)

Table 5: We selected several LLMs and produced LLM-generated text many times, and then conducted inspections to evaluate the impact of diversity on the detector.

following the dataset construction methodology used in DetectGPT and Fast-detectGPT, which are widely recognized in the literature on AI-generated text detection. This setting minimizes the influence of factors, other than model evolution, on the detection performance.

To assess the domain generalization capability of detection methods, we included the following datasets: Writing, PubMed, Community, and PeerRead. We also incorporated three different generation paradigms: continuation, question-answering, and paraphrasing. Specifically, the prompts used for generating each dataset are as follows:

- **Xsum**: "Please write an article with about 150 words starting exactly with: <prefix>"
- **Writing**: "Please write an article with about 150 words starting exactly with: <prefix>"
- **PubMed**: "Please answer the question in about 50 words: <question>"
- **SocialMedia**: "Generate text similar to the input social media text but using different words and sentence composition: <Full text>"
- **PeerRead**: "Please write a peer review with about 150 words starting exactly with: <prefix>"

F Experimental Setup Details

In this section, we provide a detailed description of the experimental setup. Specifically, for zero-shot, we follow Fast-DetectGPT (Bao et al.) and use GPT-J-6B (Wang and Komatsuzaki, 2021) and GPT-Neo-2.7B (Black et al., 2021) as the scoring models for zero-shot methods. For supervised detectors, we used the pre-trained detectors provided by the authors. All experiments are conducted on a workstation equipped with 4 NVIDIA A100 GPUs.

For the valuation metric, we measure the detection performance using two metrics: AUROC(the area under the receiver operating characteristic) and EMG (evolving model generalization). A higher AUROC value indicates better detection quality, while a higher EMG metric indicates better generalization ability across evolving LLMs.

G The explanation of EMG

In this section, we provide an explanation of EMG. A positive EMG value indicates that the detection performance improves as the LLM evolves, while negative values reflect a decline in generalization. The magnitude of the EMG value also reflects the degree of improvement or decline in generalization; smaller values indicate worse stability. EMG reflects the direction and degree of fluctuations. However, it does not directly indicate the actual detection performance regarding AUROC.

For example, as shown in Table 2, in the GPT-4o family, the EMG values for RADAR and Text-Fluoroscopy are -0.537% and -0.493%, respectively. However, as shown in Table 1, the AUROC values for RADAR range from 0.8416 to 0.8694, while those for Text-Fluoroscopy fluctuate between 0.75426 to 0.79618. Although the EMG values are similar, the AUROC differs significantly, suggesting that both methods have comparable generalization abilities in the face of evolving LLMs, but their actual detection performance varies. Therefore, we recommend evaluating both AUROC and EMG together.

H Results of Other 12 Detection Methods

In this section, we provide the detailed results of the other detection methods in Table 7, 8, 9, 10, 11, 12, 13. The results are presented with four decimal places in the tables.

Table 6: Details of evolving LLMs.

Evolving LLMs	**Version Time**	**Source**
GPT-4o	2024-05-13	gpt-4o-2024-05-13
GPT-4o	2024-08-06	gpt-4o-2024-08-06
GPT-4o	2024-11-20	gpt-4o-2024-11-20
GPT-4o	Latest	chatgpt-4o-latest
GPT-4o-mini	2024-07-18	gpt-4o-mini-2024-07-18
GPT-4	2023-06-13	gpt-4-0613
GPT-4	2023-11-06	gpt-4-1106-preview
GPT-4	2024-01-25	gpt-4-0125-preview
GPT-4	2024-04-09	gpt-4-turbo-2024-04-09
Claude-Sonnet	2024-02-29	claude-3-sonnet-20240229
Claude-Sonnet	2024-06-20	claude-3-5-sonnet-20240620
Claude-Sonnet	2024-10-22	claude-3-5-sonnet-20241022
Claude-Haiku	2024-03-07	claude-3-haiku-20240307
Claude-Haiku	2024-10-22	claude-3-5-haiku-20241022
Claude-Opus	2024-02-29	claude-3-opus-20240229
Qwen	Qwen1.5-7B	Qwen/Qwen1.5-7B-Chat
Qwen	Qwen2-7B	Qwen2-7B-Instruct
Qwen	Qwen2.5-7B	Qwen/Qwen2.5-7B-Instruct
LlaMa3	Llama-3.1-8B	meta-llama/Meta-Llama-3.1-8B-Instruct
LlaMa3	Llama-3.1-70B	meta-llama/Meta-Llama-3.1-70B-Instruct
LlaMa3	Llama-3.2-1B	meta-llama/Meta-Llama-3.2-1B-Instruct
LlaMa3	Llama-3.2-3B	meta-llama/Meta-Llama-3.2-3B-Instruct
LlaMa3	Llama-3.3-70B	meta-llama/Meta-Llama-3.3-70B-Instruct
Fine-tuning	LlaMA-2-7B-chat-hf	meta-llama/Llama-2-7b-chat-hf
Fine-tuning	Vicuna-1.5-7B	lmsys/vicuna-7b-v1.5
Fine-tuning	Wizardmath-7B	WizardLMTeam/WizardMath-7B-V1.0
Pruning	Sheared-LlaMA1.3B	princeton-nlp/Sheared-LLaMA-1.3B
Pruning	Sheared-LlaMA1.3B-pruned	princeton-nlp/Sheared-LLaMA-1.3B-Pruned
Pruning	Sheared-LlaMA2.7B-pruned	princeton-nlp/Sheared-LLaMA-2.7B
Pruning	Sheared-LlaMA2.7B	princeton-nlp/Sheared-LLaMA-2.7B-Pruned

LLMs	Version Time/ Version Name	Fast-DetectGPT						RADAR					
		Xsum	Writing	PubMed	SocialMedia	PeerRead	Avg.	Xsum	Writing	PubMed	SocialMedia	PeerRead	Avg.
Updates													
GPT-4o	2024-05-13	0.8971	0.9688	0.7346	0.5756	0.8256	0.8003	0.9945	0.8281	0.8276	0.5976	0.6781	0.7852
	2024-08-06	0.8696 (−2.75%)	0.9447 (−2.41%)	0.6992 (−3.54%)	0.5932 (+1.76%)	0.7667 (−5.89%)	0.7747 (−2.57%)	0.9934 (−0.11%)	0.8608 (+3.27%)	0.8223 (−0.53%)	0.6188 (+2.12%)	0.6856 (+0.75%)	0.7962 (+1.10%)
	2024-11-20	0.7202 (−17.6%)	0.8998 (−6.90%)	0.6969 (−3.77%)	0.5747 (−0.09%)	0.7831 (−4.25%)	0.7349 (−6.54%)	0.9858 (−0.87%)	0.7175 (−11.0%)	0.7868 (−4.08%)	0.6178 (+2.02%)	0.6634 (−1.47%)	0.7543 (−3.09%)
	Latest	0.6963 (−20.0%)	0.9075 (−6.13%)	0.7010 (−3.36%)	0.5820 (+0.64%)	0.8242 (−0.14%)	0.7422 (−5.81%)	0.9866 (−0.79%)	0.7532 (−7.49%)	0.7993 (−2.83%)	0.6422 (+4.46%)	0.6728 (−0.53%)	0.7708 (−1.44%)
GPT-4o-mini	2024-07-18	0.9062 (+0.91%)	0.9700 (+0.12%)	0.7304 (−0.42%)	0.6058 (+3.02%)	0.8042 (−2.14%)	0.8033 (+0.30%)	0.9978 (+0.33%)	0.8663 (+3.82%)	0.8208 (−0.68%)	0.6645 (+6.69%)	0.6385 (−3.96%)	0.7976 (+1.24%)
GPT-4	2023-06-13	0.9262	0.9776	0.7611	0.4937	0.7511	0.7819	0.9944	0.8725	0.8918	0.6562	0.7490	0.8328
	2023-11-06	0.9177 (−0.85%)	0.9300 (−4.76%)	0.7317 (−2.94%)	0.5480 (+5.43%)	0.8282 (+7.71%)	0.7911 (+0.92%)	0.9943 (−0.01%)	0.7876 (−8.49%)	0.8037 (−8.81%)	0.6586 (+0.24%)	0.6259 (−12.3%)	0.7740 (−5.88%)
	2024-01-25	0.8940 (−3.22%)	0.9435 (−3.41%)	0.7036 (−5.75%)	0.5360 (+4.23%)	0.8245 (+7.34%)	0.7803 (−0.16%)	0.9921 (−0.23%)	0.8002 (−7.23%)	0.8409 (−5.09%)	0.6516 (−0.46%)	0.6622 (−8.68%)	0.7894 (−4.34%)
	2024-04-09	0.8332 (−9.30%)	0.9120 (−6.56%)	0.7172 (−4.39%)	0.5764 (+8.27%)	0.7919 (+4.08%)	0.7661 (−1.58%)	0.9908 (−0.36%)	0.7880 (−8.45%)	0.8211 (−7.07%)	0.5970 (−5.92%)	0.6008 (−14.8%)	0.7595 (−7.32%)
Claude-Sonnet	2024-02-29	0.9459	0.9804	0.8150	0.5836	0.9320	0.8514	0.9700	0.8582	0.8403	0.7355	0.6106	0.8029
	2024-06-20	0.9692 (+2.33%)	0.9884 (+0.80%)	0.7718 (−4.32%)	0.6605 (+7.69%)	0.9030 (−2.90%)	0.8586 (+0.72%)	0.9861 (+1.61%)	0.8061 (−5.21%)	0.8159 (−2.44%)	0.7250 (−1.05%)	0.5549 (−5.57%)	0.7776 (−2.53%)
	2024-10-22	0.9016 (−4.43%)	0.9265 (−5.39%)	0.6760 (−13.9%)	0.5056 (−7.80%)	0.7789 (−15.3%)	0.7577 (−9.37%)	0.9548 (−1.52%)	0.7238 (−13.4%)	0.7418 (−9.85%)	0.6681 (−6.74%)	0.5468 (−6.38%)	0.7271 (−7.59%)
Claude-Haiku	2024-03-07	0.9952	0.9993	0.8614	0.7486	0.9389	0.9087	0.9972	0.9328	0.8388	0.7669	0.6745	0.8420
	2024-10-22	0.8396 (−15.5%)	0.9237 (−7.56%)	0.6439 (−21.7%)	0.3892 (−35.9%)	0.4641 (−47.4%)	0.6521 (−25.6%)	0.9965 (−0.07%)	0.8578 (−7.50%)	0.7293 (−10.9%)	0.7031 (−6.38%)	0.8217 (+14.7%)	0.8217 (−2.04%)
Claude-Opus	2024-02-29	0.9703 (−2.49%)	0.9604 (−3.89%)	0.8234 (−3.80%)	0.7159 (−3.27%)	0.8943 (−4.46%)	0.8729 (−3.58%)	0.9939 (−0.33%)	0.8727 (−6.01%)	0.8165 (−2.23%)	0.7698 (+0.29%)	0.6188 (−5.57%)	0.8143 (−2.77%)
Qwen	Qwen1.5-7B	0.9243	0.9945	0.6641	0.4897	0.9076	0.7960	0.9269	0.9034	0.7222	0.5074	0.4761	0.7072
	Qwen2-7B	0.9929 (+6.86%)	0.9974 (+0.29%)	0.7376 (+7.35%)	0.4622 (−2.75%)	0.8939 (−1.37%)	0.8168 (+2.08%)	0.8856 (−4.13%)	0.8692 (−3.42%)	0.8168 (+9.46%)	0.5209 (+1.35%)	0.5375 (+6.14%)	0.7260 (+1.88%)
	Qwen2.5-7B	0.9894 (+6.51%)	0.9955 (+0.10%)	0.7878 (+12.3%)	0.4686 (−2.11%)	0.8998 (−0.78%)	0.8282 (+3.22%)	0.9009 (−2.60%)	0.9100 (+0.66%)	0.7907 (+6.85%)	0.5205 (+1.31%)	0.5256 (+4.95%)	0.7295 (+2.23%)
LlaMA3	LlaMA-3.1-8B	0.9926	0.9839	0.8476	0.6866	0.9680	0.8957	0.9987	0.8939	0.7784	0.7232	0.6304	0.8049
	LlaMA-3.1-70B	0.9957 (+0.31%)	0.9989 (+1.50%)	0.8104 (−3.72%)	0.6693 (−1.73%)	0.9584 (−0.96%)	0.8865 (−0.92%)	0.9968 (−0.19%)	0.8807 (−1.32%)	0.7627 (−1.57%)	0.7361 (+1.29%)	0.6182 (−1.22%)	0.7989 (−0.60%)
	LlaMA-3.2-1B	0.9580 (−3.46%)	0.9948 (+1.09%)	0.9205 (+7.29%)	0.8617 (+17.5%)	0.9650 (−0.30%)	0.9400 (+4.43%)	0.9636 (−3.51%)	0.9378 (+4.39%)	0.8158 (+3.74%)	0.7512 (+2.80%)	0.7056 (+7.52%)	0.8348 (+2.99%)
	LlaMA-3.2-3B	0.9913 (−0.13%)	0.9870 (+0.31%)	0.9132 (+6.56%)	0.7698 (+8.32%)	0.9650 (−0.30%)	0.9253 (+2.95%)	0.9943 (−0.44%)	0.9210 (+2.71%)	0.8154 (+3.70%)	0.7655 (+4.23%)	0.6584 (+2.80%)	0.8309 (+2.60%)
	LlaMA-3.3-70B	0.9958 (+0.32%)	1.0000 (+1.61%)	0.7923 (−5.53%)	0.6881 (+0.15%)	0.9042 (−6.38%)	0.8761 (−1.97%)	0.9955 (−0.32%)	0.8865 (−0.74%)	0.7878 (+0.94%)	0.7308 (+0.76%)	0.6533 (+2.29%)	0.8108 (+0.59%)
Developments													
Fine-tuning	LlaMA-2-7B-chat-hf	0.9722	0.9904	0.9055	0.8963	0.9693	0.9467	0.6808	0.6330	0.4911	0.6396	0.6380	0.6165
	Vicuna-1.5-7B	0.9958 (+2.36%)	0.9956 (+0.52%)	0.8008 (−10.4%)	0.8805 (−1.58%)	0.9439 (−2.54%)	0.9233 (−2.34%)	0.9745 (+29.3%)	0.9305 (+29.7%)	0.8064 (+31.5%)	0.6490 (+0.94%)	0.6355 (−0.25%)	0.7992 (+18.2%)
	Wizardmath-7B	0.9205 (−5.17%)	0.9804 (−1.00%)	0.7354 (−17.0%)	0.8490 (−4.73%)	0.8412 (−12.8%)	0.8653 (−8.14%)	0.7476 (+6.68%)	0.7582 (+12.5%)	0.6969 (+20.5%)	0.6857 (+4.61%)	0.5447 (−9.33%)	0.6866 (+7.01%)
	Llama-2-7B-hf	0.6676 (−30.4%)	0.9023 (−8.81%)	0.4810 (−42.4%)	0.7783 (−11.8%)	0.7594 (−20.9%)	0.7177 (−22.9%)	0.9168 (+23.6%)	0.8058 (+17.2%)	0.6961 (+20.5%)	0.5787 (−6.09%)	0.4464 (−19.1%)	0.6888 (+7.23%)
Pruning	Sheared-LlaMA1.3B	0.6548 (−31.7%)	0.8888 (−10.1%)	0.4846 (−42.0%)	0.8634 (−3.29%)	0.7072 (−26.2%)	0.7198 (−22.7%)	0.7706 (+8.98%)	0.6626 (+2.96%)	0.5856 (+9.45%)	0.7255 (+8.59%)	0.5360 (−10.2%)	0.6561 (+3.96%)
	Sheared-LlaMA1.3B-pruned	0.4118 (−56.0%)	0.8077 (−18.2%)	0.3524 (−55.3%)	0.7530 (−14.3%)	0.4593 (−51.0%)	0.5568 (−38.9%)	0.9213 (+24.0%)	0.7903 (+15.7%)	0.6991 (+20.8%)	0.8669 (+22.7%)	0.6539 (+1.59%)	0.7863 (+16.9%)
	Sheared-LlaMA2.7B-pruned	0.4480 (−52.4%)	0.8672 (−12.3%)	0.3664 (−53.9%)	0.7543 (−14.2%)	0.5196 (−44.9%)	0.5911 (−35.5%)	0.8383 (+15.7%)	0.7618 (+12.8%)	0.5870 (+9.59%)	0.7831 (+14.3%)	0.5455 (−9.25%)	0.7031 (+8.66%)
	Sheared-LlaMA2.7B	0.7049 (−26.7%)	0.9154 (−7.50%)	0.5047 (−40.0%)	0.8191 (−7.72%)	0.6520 (−31.7%)	0.7192 (−22.7%)	0.7496 (+6.88%)	0.6811 (+4.81%)	0.5794 (+8.83%)	0.6743 (+3.47%)	0.4655 (−17.2%)	0.6300 (+1.35%)

Table 7: The detection performance (measured in AUROC) of Fast-DetectGPT and RADAR on EvoBench.

LLMs	Version Time/ Version Name	Text-Fluoroscopy						Imitate Before Detect					
		Xsum	Writing	PubMed	SocialMedia	PeerRead	Avg.	Xsum	Writing	PubMed	SocialMedia	PeerRead	Avg.
Updates													
GPT-4o	2024-05-13	0.9994	0.9823	0.8510	0.6870	0.8124	0.8664	0.9964	0.9835	0.7189	0.6983	0.9126	0.8619
	2024-08-06	0.9995 (+0.01%)	0.9867 (+0.44%)	0.8423 (−0.87%)	0.7072 (+2.02%)	0.8116 (−0.08%)	0.8695 (+0.30%)	0.9974 (+0.10%)	0.9836 (+0.01%)	0.7254 (+0.65%)	0.7416 (+4.33%)	0.9056 (−0.70%)	0.8707 (+0.88%)
	2024-11-20	0.9997 (+0.03%)	0.9584 (−2.39%)	0.8130 (−3.80%)	0.6877 (+0.07%)	0.7496 (−6.28%)	0.8417 (−2.47%)	0.9802 (−1.62%)	0.9480 (−3.55%)	0.6889 (−3.00%)	0.7358 (+3.75%)	0.8998 (−1.28%)	0.8505 (−1.14%)
	Latest	0.9994 (−0.00%)	0.9233 (−5.90%)	0.8020 (−4.90%)	0.7300 (+4.30%)	0.8293 (+1.69%)	0.8568 (−0.96%)	0.9838 (−1.26%)	0.9566 (−2.69%)	0.6847 (−3.42%)	0.7416 (+4.33%)	0.9121 (−0.05%)	0.8558 (−0.62%)
GPT-4o-mini	2024-07-18	1.0000 (+0.06%)	0.9894 (+0.71%)	0.8438 (−0.72%)	0.7185 (+3.15%)	0.8369 (+2.45%)	0.8777 (+1.13%)	0.9974 (+0.10%)	0.9897 (+0.62%)	0.7555 (+3.66%)	0.7887 (+9.04%)	0.9117 (−0.09%)	0.8886 (+2.67%)
GPT-4	2023-06-13	0.9997	0.9850	0.8748	0.6930	0.7289	0.8563	0.9951	0.9877	0.7094	0.6753	0.8619	0.8459
	2023-11-06	0.9995 (−0.02%)	0.9722 (−1.28%)	0.8230 (−5.18%)	0.7298 (+3.68%)	0.7368 (+0.79%)	0.8523 (−0.40%)	0.9970 (+0.19%)	0.9684 (−1.93%)	0.7907 (+8.13%)	0.7854 (+11.0%)	0.9071 (+4.52%)	0.8897 (+4.38%)
	2024-01-25	0.9998 (+0.01%)	0.9750 (−1.00%)	0.8079 (−6.69%)	0.7099 (+1.69%)	0.7397 (+1.08%)	0.8465 (−0.98%)	0.9956 (+0.05%)	0.9708 (−1.69%)	0.7433 (+3.39%)	0.7214 (+4.61%)	0.8936 (+3.17%)	0.8649 (+1.91%)
	2024-04-09	0.9998 (+0.01%)	0.9779 (−0.71%)	0.8171 (−5.77%)	0.6901 (−0.29%)	0.7671 (+3.82%)	0.8504 (−0.59%)	0.9930 (−0.21%)	0.9615 (−2.62%)	0.7136 (+0.42%)	0.7186 (+4.33%)	0.8930 (+3.11%)	0.8559 (+1.01%)
Claude-Sonnet	2024-02-29	0.9913	0.9810	0.8494	0.7582	0.7972	0.8754	0.9804	0.9850	0.7732	0.8864	0.9080	0.9066
	2024-06-20	0.9890 (−0.23%)	0.9759 (−0.51%)	0.8721 (+2.27%)	0.8188 (+6.06%)	0.8468 (+4.96%)	0.9005 (+2.51%)	0.9957 (+1.53%)	0.9883 (+0.33%)	0.8092 (+3.60%)	0.9127 (+2.63%)	0.9057 (−0.23%)	0.9223 (+1.57%)
	2024-10-22	0.9906 (−0.07%)	0.9421 (−3.89%)	0.8096 (−3.98%)	0.7843 (+2.61%)	0.8167 (+1.95%)	0.8687 (−0.68%)	0.9916 (+1.12%)	0.9564 (−2.86%)	0.7138 (−5.94%)	0.7866 (−9.98%)	0.8437 (−6.43%)	0.8584 (−4.82%)
Claude-Haiku	2024-03-07	0.9998	0.9856	0.8482	0.8182	0.7997	0.8903	0.9996	0.9981	0.8428	0.9192	0.9331	0.9386
	2024-10-22	0.9999 (+0.01%)	0.9650 (−2.06%)	0.8436 (−0.46%)	0.7944 (−2.38%)	0.8591 (+5.94%)	0.8924 (+0.21%)	0.9981 (−0.15%)	0.9796 (−1.85%)	0.7884 (−5.44%)	0.8375 (−8.17%)	0.7420 (−19.1%)	0.8691 (−6.94%)
Claude-Opus	2024-02-29	0.9968 (−0.30%)	0.9801 (−0.55%)	0.8403 (−0.79%)	0.7974 (−2.08%)	0.8140 (+1.43%)	0.8857 (−0.46%)	0.9899 (−0.97%)	0.9803 (−1.78%)	0.8080 (−3.48%)	0.9116 (−0.76%)	0.8950 (−3.81%)	0.9170 (−2.16%)
Qwen	Qwen1.5-7B	0.9017	0.9581	0.7955	0.5083	0.5593	0.7446	0.9493	0.9822	0.7517	0.4690	0.7928	0.7890
	Qwen2-7B	0.6711 (−23.0%)	0.9378 (−2.03%)	0.7617 (−3.38%)	0.4819 (−2.64%)	0.5746 (+1.53%)	0.6854 (−5.92%)	0.9765 (+2.72%)	0.9883 (+0.61%)	0.7018 (−4.99%)	0.4779 (+0.89%)	0.7784 (−1.44%)	0.7846 (−0.44%)
	Qwen2.5-7B	0.7734 (−12.8%)	0.9457 (−1.24%)	0.7208 (−7.47%)	0.4782 (−3.01%)	0.5731 (+1.38%)	0.6982 (−4.63%)	0.9676 (+1.83%)	0.9856 (+0.34%)	0.7594 (+0.77%)	0.4768 (+0.78%)	0.7801 (−1.27%)	0.7939 (+0.49%)
LlaMA3	LlaMA-3.1-8B	0.9997	0.9795	0.8108	0.7554	0.6524	0.8396	0.9988	0.9896	0.8630	0.8855	0.9040	0.9282
	LlaMA-3.1-70B	0.9999 (+0.02%)	0.9908 (+1.13%)	0.7969 (−1.39%)	0.7700 (+1.46%)	0.6179 (−3.45%)	0.8351 (−0.45%)	0.9981 (−0.07%)	0.9957 (+0.61%)	0.7594 (−10.3%)	0.9103 (+2.48%)	0.8935 (−1.05%)	0.9114 (−1.68%)
	LlaMA-3.2-1B	0.9838 (−1.59%)	0.9844 (+0.49%)	0.8541 (+4.33%)	0.8428 (+8.74%)	0.6758 (+2.34%)	0.8682 (+2.86%)	0.9915 (−0.73%)	0.9963 (+0.67%)	0.8981 (+3.51%)	0.9476 (+6.21%)	0.9146 (+1.06%)	0.9496 (+2.14%)
	LlaMA-3.2-3B	0.9994 (−0.03%)	0.9779 (−0.16%)	0.8383 (+2.75%)	0.8077 (+5.23%)	0.6673 (+1.49%)	0.8581 (+1.86%)	0.9992 (+0.04%)	0.9952 (+0.56%)	0.8818 (+1.88%)	0.9579 (+7.24%)	0.9192 (+1.52%)	0.9507 (+2.25%)
	LlaMA-3.3-70B	1.0000 (+0.03%)	0.9883 (+0.88%)	0.8240 (+1.32%)	0.7978 (+4.24%)	0.7587 (+10.6%)	0.8738 (+3.42%)	0.9969 (−0.19%)	0.9954 (+0.58%)	0.7246 (−13.8%)	0.9142 (+2.87%)	0.8629 (−4.11%)	0.8988 (−2.94%)
Developments													
Fine-tuning	LlaMA-2-7B-chat-hf	0.8320	0.8940	0.7480	0.8427	0.7311	0.8096	0.7310	0.7203	0.5381	0.8248	0.9436	0.7516
	Vicuna-1.5-7B	0.9138 (+8.18%)	0.9150 (+2.10%)	0.6650 (−8.30%)	0.6332 (−20.9%)	0.6375 (−9.36%)	0.7529 (−5.67%)	0.9974 (+26.6%)	0.9952 (+27.4%)	0.7742 (+23.6%)	0.7408 (−8.40%)	0.8704 (−7.32%)	0.8756 (+12.4%)
	Wizardmath-7B	0.5421 (−28.9%)	0.7817 (−11.2%)	0.5774 (−17.0%)	0.6552 (−18.7%)	0.6048 (−12.6%)	0.6322 (−17.7%)	0.8829 (+15.1%)	0.9255 (+20.5%)	0.7119 (+17.3%)	0.7522 (−7.26%)	0.7878 (−15.5%)	0.8121 (+6.05%)
	Llama-2-7B-hf	0.4764 (−35.5%)	0.6419 (−25.2%)	0.4751 (−27.2%)	0.7159 (−12.6%)	0.5464 (−18.4%)	0.5711 (−23.8%)	0.9748 (+24.3%)	0.9700 (+24.9%)	0.8896 (+35.1%)	0.7722 (−5.26%)	0.7270 (−21.6%)	0.8667 (+11.5%)
Pruning	Sheared-LlaMA1.3B	0.5258 (−30.6%)	0.6373 (−25.6%)	0.5212 (−22.6%)	0.8011 (−4.16%)	0.5287 (−20.2%)	0.6028 (−20.6%)	0.7478 (+1.68%)	0.7437 (+2.34%)	0.5591 (+2.10%)	0.7700 (−5.48%)	0.7246 (−21.9%)	0.7090 (−4.25%)
	Sheared-LlaMA1.3B-pruned	0.8359 (+0.39%)	0.9438 (+4.98%)	0.6540 (−9.40%)	0.7971 (−4.56%)	0.6440 (−8.71%)	0.7750 (−3.46%)	0.6762 (−5.48%)	0.5737 (−14.6%)	0.4265 (−11.1%)	0.6502 (−17.4%)	0.5600 (−38.3%)	0.5773 (−17.4%)
	Sheared-LlaMA2.7B-pruned	0.4902 (−34.1%)	0.6440 (−25.0%)	0.5300 (−21.8%)	0.6716 (−17.1%)	0.5846 (−14.6%)	0.5841 (−22.5%)	0.7472 (+1.62%)	0.7484 (+2.81%)	0.5405 (+0.24%)	0.7056 (−11.9%)	0.6300 (−31.3%)	0.6743 (−7.72%)
	Sheared-LlaMA2.7B	0.5864 (−24.5%)	0.6454 (−24.8%)	0.5215 (−22.6%)	0.6637 (−17.9%)	0.5626 (−16.8%)	0.5959 (−21.3%)	0.7030 (−2.80%)	0.6804 (−3.99%)	0.4982 (−3.99%)	0.6789 (−14.5%)	0.6021 (−34.1%)	0.6325 (−11.9%)

Table 8: The detection performance (measured in AUROC) of Text-Fluoroscopy and Imitate Before Detect on EvoBench.

LLMs	Version Time/ Version Name	RoBERTa-base						RoBERTa-large					
		Xsum	Writing	PubMed	SocialMedia	PeerRead	Avg.	Xsum	Writing	PubMed	SocialMedia	PeerRead	Avg.
Updates													
GPT-4o	2024-05-13	0.6696	0.4668	0.4680	0.4946	0.5584	0.5315	0.6124	0.3163	0.5102	0.5804	0.6031	0.5245
	2024-08-06	0.5975 (−7.21%)	0.4108 (−5.60%)	0.4644 (−0.36%)	0.5487 (+5.41%)	0.5425 (−1.59%)	0.5128 (−1.87%)	0.6478 (+3.54%)	0.3191 (+0.28%)	0.5079 (−0.23%)	0.5897 (+0.93%)	0.5646 (−3.85%)	0.5258 (+0.13%)
	2024-11-20	0.5858 (−8.38%)	0.3992 (−6.76%)	0.4331 (−3.49%)	0.4644 (−3.02%)	0.5061 (−5.23%)	0.4777 (−5.38%)	0.4226 (−18.9%)	0.2562 (−6.01%)	0.4536 (−5.66%)	0.5414 (−3.90%)	0.5402 (−6.29%)	0.4428 (−8.17%)
	Latest	0.5417 (−12.7%)	0.3969 (−6.99%)	0.4386 (−2.94%)	0.5256 (+3.10%)	0.5269 (−3.15%)	0.4859 (−4.55%)	0.4151 (−19.7%)	0.2583 (−5.80%)	0.4599 (−5.03%)	0.6150 (+3.46%)	0.5889 (−1.42%)	0.4674 (−5.70%)
GPT-4o-mini	2024-07-18	0.7376 (+6.80%)	0.4991 (+3.23%)	0.4542 (−1.38%)	0.5608 (+6.62%)	0.5314 (−2.70%)	0.5566 (+2.51%)	0.7445 (+13.2%)	0.4498 (+13.3%)	0.5256 (+1.54%)	0.6047 (+2.43%)	0.5887 (−1.44%)	0.5827 (+5.82%)
GPT-4	2023-06-13	0.7349	0.5198	0.5286	0.5664	0.5205	0.5740	0.7368	0.3972	0.5917	0.6340	0.5932	0.5906
	2023-11-06	0.6978 (−3.71%)	0.4235 (−9.63%)	0.4483 (−8.03%)	0.5488 (−1.76%)	0.5584 (+3.79%)	0.5354 (−3.87%)	0.6471 (−8.97%)	0.3444 (−5.28%)	0.5056 (−8.61%)	0.5924 (−4.16%)	0.6205 (+2.73%)	0.5420 (−4.86%)
	2024-01-25	0.6728 (−6.21%)	0.4948 (−2.50%)	0.4545 (−7.41%)	0.5452 (−2.12%)	0.5484 (+2.79%)	0.5431 (−3.09%)	0.6634 (−7.34%)	0.3589 (−3.83%)	0.5221 (−6.96%)	0.5901 (−4.39%)	0.6137 (+2.05%)	0.5496 (−4.09%)
	2024-04-09	0.5949 (−14.0%)	0.3947 (−12.5%)	0.4817 (−4.69%)	0.5029 (−6.35%)	0.5430 (+2.25%)	0.5034 (−7.06%)	0.6030 (−13.3%)	0.2916 (−10.5%)	0.5179 (−7.38%)	0.5550 (−7.90%)	0.5820 (−1.12%)	0.5099 (−8.07%)
Claude-Sonnet	2024-02-29	0.7459	0.5915	0.5312	0.6558	0.5397	0.6128	0.6757	0.4837	0.5679	0.7196	0.6077	0.6109
	2024-06-20	0.7550 (+0.91%)	0.5054 (−8.61%)	0.4776 (−5.36%)	0.5935 (−6.23%)	0.5355 (−0.42%)	0.5734 (−3.94%)	0.6328 (−4.29%)	0.3337 (−15.0%)	0.5254 (−4.25%)	0.6633 (−5.63%)	0.5861 (−2.16%)	0.5483 (−6.27%)
	2024-10-22	0.6080 (−13.7%)	0.3702 (−22.1%)	0.3741 (−15.7%)	0.5283 (−12.7%)	0.5082 (−3.15%)	0.4778 (−13.5%)	0.5696 (−10.6%)	0.2660 (−21.7%)	0.4148 (−15.3%)	0.5533 (−16.6%)	0.5147 (−9.30%)	0.4637 (−14.7%)
Claude-Haiku	2024-03-07	0.9427	0.8251	0.5400	0.6579	0.6523	0.7236	0.9162	0.6806	0.6079	0.7053	0.6716	0.7163
	2024-10-22	0.6408 (−30.1%)	0.3626 (−46.2%)	0.3556 (−18.4%)	0.5740 (−8.39%)	0.5146 (−13.7%)	0.4895 (−23.4%)	0.6422 (−27.4%)	0.2771 (−40.3%)	0.4116 (−19.6%)	0.6527 (−5.26%)	0.5997 (−7.19%)	0.5167 (−19.9%)
Claude-Opus	2024-02-29	0.8337 (−10.9%)	0.6305 (−19.4%)	0.4875 (−5.25%)	0.5825 (−7.54%)	0.5565 (−9.58%)	0.6181 (−10.5%)	0.7674 (−14.8%)	0.4963 (−18.4%)	0.5180 (−8.99%)	0.6711 (−3.42%)	0.5954 (−7.62%)	0.6096 (−10.6%)
Qwen	Qwen1.5-7B	0.8357	0.8064	0.6017	0.4882	0.5339	0.6532	0.8891	0.7393	0.6413	0.5203	0.5213	0.6623
	Qwen2-7B	0.9770 (+14.1%)	0.9259 (+11.9%)	0.6218 (+2.01%)	0.4980 (+0.98%)	0.5404 (+0.65%)	0.7126 (+5.94%)	0.9624 (+7.33%)	0.8716 (+13.2%)	0.6588 (+1.75%)	0.4953 (−2.50%)	0.5440 (+2.27%)	0.7064 (+4.42%)
	Qwen2.5-7B	0.9776 (+14.1%)	0.8910 (+8.46%)	0.6249 (+2.32%)	0.5027 (+1.45%)	0.5592 (+2.53%)	0.7111 (+5.79%)	0.9645 (+7.54%)	0.8508 (+11.1%)	0.6841 (+4.28%)	0.5017 (−1.86%)	0.5552 (+3.39%)	0.7113 (+4.90%)
LlaMA3	LlaMA-3.1-8B	0.9590	0.8822	0.5254	0.5528	0.6138	0.7066	0.9344	0.7625	0.5748	0.5896	0.6216	0.6966
	LlaMA-3.1-70B	0.9676 (+0.86%)	0.8311 (−5.11%)	0.4566 (−6.88%)	0.5980 (+4.52%)	0.5769 (−3.69%)	0.6860 (−2.06%)	0.9080 (−2.64%)	0.7065 (−5.60%)	0.5002 (−7.46%)	0.6352 (+4.56%)	0.5824 (−3.92%)	0.6665 (−3.01%)
	LlaMA-3.2-1B	0.9983 (+3.93%)	0.9664 (+8.42%)	0.6148 (+8.94%)	0.6376 (+8.48%)	0.6326 (+1.88%)	0.7699 (+6.33%)	0.9946 (+6.02%)	0.9399 (+17.7%)	0.7275 (+15.2%)	0.7210 (+13.1%)	0.6772 (+5.56%)	0.8120 (+11.5%)
	LlaMA-3.2-3B	0.9748 (+1.58%)	0.9216 (+3.94%)	0.5327 (+0.73%)	0.6107 (+5.79%)	0.6326 (+1.88%)	0.7345 (+2.78%)	0.9736 (+3.92%)	0.8459 (+8.34%)	0.6316 (+5.68%)	0.6768 (+8.72%)	0.6772 (+5.56%)	0.7610 (+6.44%)
	LlaMA-3.3-70B	0.9534 (−0.56%)	0.8542 (−2.80%)	0.4533 (−7.21%)	0.6084 (+5.56%)	0.6123 (−0.15%)	0.6963 (−1.03%)	0.9315 (−0.29%)	0.7169 (−4.56%)	0.5106 (−6.42%)	0.6604 (+7.08%)	0.6438 (+2.22%)	0.6926 (−0.39%)
Developments													
Fine-tuning	LlaMA-2-7B-chat-hf	0.9314	0.8601	0.6764	0.6858	0.6604	0.7628	0.9249	0.8129	0.7219	0.7028	0.6582	0.7641
	Vicuna-1.5-7B	0.9606 (+2.92%)	0.9720 (+11.1%)	0.6742 (−0.22%)	0.5445 (−14.1%)	0.5766 (−8.38%)	0.7456 (−1.72%)	0.9647 (+3.98%)	0.9424 (+12.9%)	0.6928 (−2.91%)	0.5659 (−13.6%)	0.6096 (−4.86%)	0.7551 (−0.91%)
	Wizardmath-7B	0.8680 (−6.34%)	0.8914 (+3.13%)	0.6902 (+1.38%)	0.6733 (−1.25%)	0.5380 (−12.2%)	0.7322 (−3.06%)	0.8845 (−4.04%)	0.8839 (+7.10%)	0.7396 (+1.77%)	0.7204 (+1.76%)	0.5649 (−9.33%)	0.7587 (−0.55%)
	Llama-2-7B-hf	0.8712 (−6.02%)	0.8686 (+0.85%)	0.6093 (−6.71%)	0.7207 (+3.49%)	0.6179 (−4.25%)	0.7375 (−2.53%)	0.9268 (+0.19%)	0.9083 (+9.54%)	0.6849 (−3.70%)	0.7586 (+5.58%)	0.6668 (+0.86%)	0.7891 (+2.49%)
Pruning	Sheared-LlaMA1.3B	0.9268 (−0.46%)	0.9089 (+4.88%)	0.6990 (+2.26%)	0.8616 (+17.5%)	0.7411 (+8.07%)	0.8275 (+6.47%)	0.9864 (+6.15%)	0.9868 (+17.3%)	0.7582 (+3.63%)	0.9155 (+21.2%)	0.7953 (+13.7%)	0.8884 (+12.4%)
	Sheared-LlaMA1.3B-pruned	0.9966 (+6.52%)	0.9924 (+13.2%)	0.8346 (+15.8%)	0.9131 (+22.7%)	0.7516 (+9.12%)	0.8977 (+13.4%)	1.0000 (+7.51%)	1.0000 (+18.7%)	0.9345 (+21.2%)	0.9827 (+27.9%)	0.8686 (+21.0%)	0.9572 (+19.3%)
	Sheared-LlaMA2.7B-pruned	0.8837 (−4.77%)	0.8977 (+3.76%)	0.6818 (+0.54%)	0.7756 (+8.98%)	0.5835 (−7.69%)	0.7645 (+0.16%)	0.9693 (+4.44%)	0.9601 (+14.7%)	0.7168 (−0.51%)	0.8469 (+14.4%)	0.6222 (−3.60%)	0.8231 (+5.89%)
	Sheared-LlaMA2.7B	0.9767 (+4.53%)	0.9649 (+10.4%)	0.7646 (+8.82%)	0.8799 (+19.4%)	0.6624 (+0.20%)	0.8497 (+8.69%)	0.9996 (+7.47%)	0.9999 (+18.7%)	0.8647 (+14.2%)	0.9480 (+24.5%)	0.7589 (+10.0%)	0.9142 (+15.0%)

Table 9: The detection performance (measured in AUROC) of RoBERTa-base and RoBERTa-large on EvoBench.

LLMs	Version Time/ Version Name	Likelihood						Rank					
		Xsum	Writing	PubMed	SocialMedia	PeerRead	Avg.	Xsum	Writing	PubMed	SocialMedia	PeerRead	Avg.
Updates													
GPT-4o	2024-05-13	0.7820	0.8560	0.8033	0.5452	0.7356	0.7444	0.6596	0.7265	0.5969	0.5359	0.6840	0.6406
	2024-08-06	0.6980 (−8.40%)	0.8023 (−5.37%)	0.7717 (−3.16%)	0.5578 (+1.26%)	0.6823 (−5.33%)	0.7024 (−4.20%)	0.6751 (+1.55%)	0.7358 (+0.93%)	0.5920 (−0.49%)	0.5422 (+0.63%)	0.6643 (−1.97%)	0.6419 (+0.13%)
	2024-11-20	0.8764 (+9.44%)	0.8361 (−1.99%)	0.7752 (−2.81%)	0.5889 (+4.37%)	0.6910 (−4.46%)	0.7535 (+0.91%)	0.4497 (−20.9%)	0.7078 (−1.87%)	0.6022 (+0.53%)	0.5432 (+0.73%)	0.6640 (−2.00%)	0.5934 (−4.72%)
	Latest	0.8708 (+8.88%)	0.8846 (+2.86%)	0.7745 (−2.88%)	0.5498 (+0.46%)	0.7183 (−1.73%)	0.7596 (+1.52%)	0.4356 (−22.4%)	0.6846 (−4.19%)	0.5960 (−0.09%)	0.5584 (+2.25%)	0.6768 (−0.72%)	0.5903 (−5.03%)
GPT-4o-mini	2024-07-18	0.7304 (−5.16%)	0.8513 (−0.47%)	0.7716 (−3.17%)	0.5846 (+3.94%)	0.6941 (−4.15%)	0.7264 (−1.80%)	0.7024 (+4.28%)	0.7429 (+1.64%)	0.5981 (+0.12%)	0.5400 (+0.41%)	0.6635 (−2.05%)	0.6494 (+0.88%)
GPT-4	2023-06-13	0.7840	0.8392	0.8132	0.4396	0.6313	0.7015	0.6759	0.7207	0.6013	0.4415	0.6378	0.6154
	2023-11-06	0.7546 (−2.94%)	0.7389 (−10.0%)	0.7907 (−2.25%)	0.4666 (+2.70%)	0.7287 (+9.74%)	0.6959 (−0.56%)	0.6718 (−0.41%)	0.6640 (−5.67%)	0.6010 (−0.03%)	0.4851 (+4.36%)	0.6710 (+3.32%)	0.6186 (+0.31%)
	2024-01-25	0.6970 (−8.70%)	0.8465 (+0.73%)	0.7933 (−1.99%)	0.4902 (+5.06%)	0.7316 (+10.0%)	0.7117 (+1.03%)	0.6644 (−1.15%)	0.7181 (−0.26%)	0.5956 (−0.57%)	0.5064 (+6.49%)	0.6767 (+3.89%)	0.6322 (+1.68%)
	2024-04-09	0.6850 (−9.90%)	0.7897 (−4.95%)	0.7978 (−1.54%)	0.5264 (+8.68%)	0.6777 (+4.64%)	0.6953 (−0.61%)	0.6795 (+0.36%)	0.7041 (−1.66%)	0.6011 (−0.02%)	0.4868 (+4.53%)	0.6676 (+2.98%)	0.6278 (+1.24%)
Claude-Sonnet	2024-02-29	0.8796	0.9386	0.8372	0.4578	0.8621	0.7951	0.7069	0.7629	0.6062	0.4769	0.7132	0.6532
	2024-06-20	0.9123 (+3.27%)	0.9338 (−0.48%)	0.7980 (−3.92%)	0.5643 (+10.6%)	0.7927 (−6.94%)	0.8002 (+0.52%)	0.6962 (−1.07%)	0.7419 (−2.10%)	0.6055 (−0.07%)	0.5523 (+7.54%)	0.6988 (−1.44%)	0.6589 (+0.57%)
	2024-10-22	0.8453 (−3.43%)	0.8024 (−13.6%)	0.7476 (−8.96%)	0.4584 (+0.06%)	0.6976 (−16.4%)	0.7103 (−8.48%)	0.6976 (−0.93%)	0.6954 (−6.75%)	0.6016 (−0.46%)	0.5305 (+5.36%)	0.6944 (−1.88%)	0.6439 (−0.93%)
Claude-Haiku	2024-03-07	0.9611	0.9821	0.8613	0.6133	0.8746	0.8585	0.7340	0.7948	0.6113	0.5825	0.7165	0.6878
	2024-10-22	0.7885 (−17.2%)	0.8390 (−14.3%)	0.6958 (−16.5%)	0.3080 (−30.5%)	0.9291 (+5.45%)	0.7121 (−14.6%)	0.4851 (−24.8%)	0.5726 (−22.2%)	0.5940 (−1.73%)	0.4323 (−15.0%)	0.1547 (−56.1%)	0.4477 (−24.0%)
Claude-Opus	2024-02-29	0.9288 (−3.23%)	0.9581 (−2.40%)	0.8415 (−1.98%)	0.6053 (−0.80%)	0.8140 (−6.06%)	0.8295 (−2.89%)	0.6964 (−3.76%)	0.7513 (−4.35%)	0.6074 (−0.39%)	0.5458 (−3.67%)	0.7006 (−1.59%)	0.6603 (−2.75%)
Qwen	Qwen1.5-7B	0.9613	0.9724	0.7250	0.5598	0.8303	0.8098	0.5473	0.7721	0.4532	0.4838	0.6862	0.5885
	Qwen2-7B	0.9897 (+2.84%)	0.9881 (+1.57%)	0.6697 (−5.53%)	0.4810 (−7.88%)	0.8378 (+0.75%)	0.7933 (−1.65%)	0.7598 (+21.2%)	0.8230 (+5.09%)	0.4206 (−3.26%)	0.4603 (−2.35%)	0.6856 (−0.06%)	0.6299 (+4.13%)
	Qwen2.5-7B	0.9834 (+2.21%)	0.9905 (+1.81%)	0.7348 (+0.98%)	0.4825 (−7.73%)	0.8221 (−0.82%)	0.8027 (−0.71%)	0.7359 (+18.8%)	0.8209 (+4.88%)	0.4551 (+0.19%)	0.4652 (−1.86%)	0.6843 (−0.19%)	0.6323 (+4.38%)
LlaMA3	LlaMA-3.1-8B	0.9568	0.9577	0.8251	0.5852	0.9287	0.8507	0.7410	0.7805	0.6072	0.5506	0.7351	0.6829
	LlaMA-3.1-70B	0.9377 (−1.91%)	0.9754 (+1.77%)	0.7928 (−3.23%)	0.5636 (−2.16%)	0.9168 (−1.19%)	0.8373 (−1.34%)	0.7304 (−1.06%)	0.7894 (+0.89%)	0.5918 (−1.54%)	0.5403 (−1.03%)	0.7212 (−1.39%)	0.6746 (−0.83%)
	LlaMA-3.2-1B	0.9406 (−1.62%)	0.9839 (+2.62%)	0.8842 (+5.91%)	0.7317 (+14.6%)	0.9244 (−0.43%)	0.8930 (+4.23%)	0.7616 (+2.06%)	0.8172 (+3.67%)	0.6192 (+1.20%)	0.6720 (+12.1%)	0.7349 (−0.02%)	0.7210 (+3.81%)
	LlaMA-3.2-3B	0.9766 (+1.98%)	0.9729 (+1.52%)	0.8718 (+4.67%)	0.6604 (+7.52%)	0.9244 (−0.43%)	0.8812 (+3.05%)	0.7503 (+0.93%)	0.8090 (+2.85%)	0.6142 (+0.70%)	0.6173 (+6.67%)	0.7349 (−0.02%)	0.7051 (+2.23%)
	LlaMA-3.3-70B	0.9324 (−2.44%)	0.9741 (+1.64%)	0.7772 (−4.79%)	0.5779 (−0.73%)	0.7860 (−14.2%)	0.8095 (−4.12%)	0.7344 (−0.66%)	0.7873 (+0.68%)	0.5992 (−0.80%)	0.5592 (+0.86%)	0.6827 (−5.24%)	0.6726 (−1.03%)
Developments													
Fine-tuning	LlaMA-2-7B-chat-hf	0.9604	0.9640	0.8568	0.8764	0.9433	0.9202	0.7120	0.7754	0.5840	0.6566	0.7868	0.7030
	Vicuna-1.5-7B	0.9961 (+3.57%)	0.9945 (+3.05%)	0.7624 (−9.44%)	0.7932 (−8.32%)	0.8948 (−4.85%)	0.8882 (−3.20%)	0.5429 (−16.9%)	0.7583 (−1.71%)	0.4262 (−15.7%)	0.4776 (−17.9%)	0.6820 (−10.4%)	0.5774 (−12.5%)
	Wizardmath-7B	0.8367 (−12.3%)	0.9352 (−2.88%)	0.6785 (−17.8%)	0.7574 (−11.9%)	0.8253 (−11.8%)	0.8066 (−11.3%)	0.6458 (−6.62%)	0.7476 (−2.78%)	0.4249 (−15.9%)	0.6588 (+0.22%)	0.6739 (−11.2%)	0.6302 (−7.28%)
	Llama-2-7B-hf	0.6500 (−31.0%)	0.8027 (−16.1%)	0.5198 (−33.7%)	0.7768 (−9.96%)	0.7620 (−18.1%)	0.7023 (−21.7%)	0.6650 (−4.70%)	0.7413 (−3.41%)	0.4683 (−11.5%)	0.7340 (+7.74%)	0.7706 (−1.62%)	0.6758 (−2.71%)
Pruning	Sheared-LlaMA1.3B	0.5791 (−38.1%)	0.7824 (−18.1%)	0.4899 (−36.6%)	0.7466 (−12.9%)	0.6532 (−29.0%)	0.6502 (−26.9%)	0.6724 (−3.96%)	0.7359 (−3.95%)	0.5050 (−7.90%)	0.7945 (+13.7%)	0.7018 (−8.50%)	0.6819 (−2.10%)
	Sheared-LlaMA1.3B-pruned	0.2053 (−75.5%)	0.6235 (−34.0%)	0.3066 (−55.0%)	0.5647 (−31.1%)	0.4081 (−53.5%)	0.4216 (−49.8%)	0.5174 (−19.4%)	0.7134 (−6.20%)	0.4448 (−13.9%)	0.6632 (+0.66%)	0.5871 (−19.9%)	0.5852 (−11.7%)
	Sheared-LlaMA2.7B-pruned	0.6160 (−34.4%)	0.8162 (−14.7%)	0.4761 (−38.0%)	0.7341 (−14.2%)	0.6598 (−28.3%)	0.6604 (−25.9%)	0.6412 (−7.08%)	0.7476 (−2.78%)	0.5184 (−6.56%)	0.6959 (+3.93%)	0.6672 (−11.9%)	0.6541 (−4.89%)
	Sheared-LlaMA2.7B	0.3425 (−61.7%)	0.6982 (−26.5%)	0.3254 (−53.1%)	0.6110 (−26.5%)	0.4911 (−45.2%)	0.4936 (−42.6%)	0.5199 (−19.2%)	0.7284 (−4.70%)	0.4564 (−12.7%)	0.6521 (−0.45%)	0.6174 (−16.9%)	0.5948 (−10.8%)

Table 10: The detection performance (measured in AUROC) of Likelihood and Rank on EvoBench.

LLMs	Version Time/ Version Name	LogRank						Entropy					
		Xsum	Writing	PubMed	SocialMedia	PeerRead	Avg.	Xsum	Writing	PubMed	SocialMedia	PeerRead	Avg.
Updates													
GPT-4o	2024-05-13	0.7804	0.8312	0.7888	0.5328	0.7234	0.7313	0.4602	0.3304	0.3091	0.5161	0.5264	0.4284
	2024-08-06	0.6969 (−8.35%)	0.7743 (−5.69%)	0.7607 (−2.81%)	0.5484 (+1.56%)	0.6670 (−5.64%)	0.6895 (−4.19%)	0.5289 (+6.87%)	0.3723 (+4.19%)	0.3269 (+1.78%)	0.4856 (−3.05%)	0.5467 (+2.03%)	0.4521 (+2.36%)
	2024-11-20	0.8408 (+6.04%)	0.8010 (−3.02%)	0.7668 (−2.20%)	0.5725 (+3.97%)	0.6758 (−4.76%)	0.7314 (+0.01%)	0.1507 (−30.9%)	0.2264 (−10.4%)	0.3105 (+0.14%)	0.4423 (−7.38%)	0.5325 (+0.61%)	0.3325 (−9.60%)
	Latest	0.8299 (+4.95%)	0.8481 (+1.69%)	0.7605 (−2.83%)	0.5412 (+0.84%)	0.6972 (−2.62%)	0.7354 (+0.41%)	0.1458 (−31.4%)	0.1671 (−16.3%)	0.3216 (+1.25%)	0.5112 (−0.49%)	0.5505 (+2.41%)	0.3392 (−8.92%)
GPT-4o-mini	2024-07-18	0.7438 (−3.66%)	0.8219 (−0.93%)	0.7630 (−2.58%)	0.5720 (+3.92%)	0.6765 (−4.69%)	0.7154 (−1.59%)	0.5408 (+8.06%)	0.3355 (+0.51%)	0.3492 (+4.01%)	0.4825 (−3.36%)	0.5636 (+3.72%)	0.4543 (+2.59%)
GPT-4	2023-06-13	0.7825	0.8139	0.8048	0.4329	0.6275	0.6923	0.4676	0.3732	0.3184	0.5743	0.5788	0.4625
	2023-11-06	0.7519 (−3.06%)	0.7061 (−10.7%)	0.7854 (−1.94%)	0.4642 (+3.13%)	0.7200 (+9.25%)	0.6855 (−0.68%)	0.5541 (+8.65%)	0.4646 (+9.14%)	0.3124 (−0.60%)	0.5726 (−0.17%)	0.5556 (−2.32%)	0.4919 (+2.94%)
	2024-01-25	0.7042 (−7.83%)	0.8260 (+1.21%)	0.7808 (−2.40%)	0.4832 (+5.03%)	0.7229 (+9.54%)	0.7034 (+1.11%)	0.5677 (+10.0%)	0.3238 (−4.94%)	0.2975 (−2.09%)	0.5511 (−2.32%)	0.5346 (−4.42%)	0.4549 (−0.75%)
	2024-04-09	0.6912 (−9.13%)	0.7563 (−5.76%)	0.7851 (−1.97%)	0.5158 (+8.29%)	0.6756 (+4.81%)	0.6848 (−0.75%)	0.5367 (+6.91%)	0.3414 (−3.18%)	0.3041 (−1.43%)	0.5134 (−6.09%)	0.5709 (−0.79%)	0.4533 (−0.92%)
Claude-Sonnet	2024-02-29	0.8854	0.9281	0.8322	0.4366	0.8678	0.7900	0.3882	0.2235	0.2986	0.5934	0.5103	0.4028
	2024-06-20	0.9047 (+1.93%)	0.9156 (−1.25%)	0.7927 (−3.95%)	0.5397 (+10.3%)	0.7974 (−7.04%)	0.7900 (−0.00%)	0.3372 (−5.10%)	0.2327 (+0.92%)	0.3312 (+3.26%)	0.5531 (−4.03%)	0.5516 (+4.13%)	0.4012 (−0.16%)
	2024-10-22	0.8404 (−4.50%)	0.7813 (−14.6%)	0.7426 (−8.96%)	0.4459 (+0.93%)	0.6936 (−17.4%)	0.7008 (−8.93%)	0.3324 (−5.58%)	0.3329 (+10.9%)	0.3414 (+4.28%)	0.5526 (−4.08%)	0.5236 (+1.33%)	0.4166 (+1.38%)
Claude-Haiku	2024-03-07	0.9675	0.9781	0.8583	0.5935	0.8824	0.8560	0.4216	0.2080	0.3073	0.5430	0.5176	0.3995
	2024-10-22	0.7687 (−19.8%)	0.7928 (−18.5%)	0.6859 (−17.2%)	0.3033 (−29.0%)	0.8854 (+0.30%)	0.6872 (−16.8%)	0.3788 (−4.28%)	0.2764 (+6.84%)	0.4005 (+9.32%)	0.6444 (+10.1%)	0.0585 (−45.9%)	0.3517 (−4.78%)
Claude-Opus	2024-02-29	0.9286 (−3.89%)	0.9469 (−3.12%)	0.8331 (−2.52%)	0.5850 (−0.85%)	0.8214 (−6.10%)	0.8230 (−3.30%)	0.3669 (−5.47%)	0.1696 (−3.84%)	0.3086 (+0.13%)	0.5333 (−0.97%)	0.5123 (−0.53%)	0.3781 (−2.14%)
Qwen	Qwen1.5-7B	0.9631	0.9723	0.7135	0.5583	0.8379	0.8090	0.2406	0.2085	0.3407	0.4227	0.4648	0.3355
	Qwen2-7B	0.9931 (+3.00%)	0.9886 (+1.63%)	0.6747 (−3.88%)	0.4851 (−7.32%)	0.8492 (+1.13%)	0.7981 (−1.09%)	0.2904 (+4.98%)	0.1665 (−4.20%)	0.4696 (+12.8%)	0.5022 (+7.95%)	0.4532 (−1.16%)	0.3764 (+4.09%)
	Qwen2.5-7B	0.9892 (+2.61%)	0.9905 (+1.82%)	0.7334 (+1.99%)	0.4858 (−7.25%)	0.8395 (+0.16%)	0.8077 (−0.13%)	0.3210 (+8.04%)	0.1523 (−5.62%)	0.4165 (+7.58%)	0.5007 (+7.80%)	0.4962 (+3.14%)	0.3773 (+4.19%)
LlaMA3	LlaMA-3.1-8B	0.9625	0.9513	0.8099	0.5759	0.9393	0.8478	0.4501	0.2765	0.3204	0.5331	0.4246	0.4009
	LlaMA-3.1-70B	0.9498 (−1.27%)	0.9726 (+2.13%)	0.7823 (−2.76%)	0.5460 (−2.99%)	0.9267 (−1.26%)	0.8355 (−1.23%)	0.5097 (+5.96%)	0.2724 (−0.41%)	0.3192 (−0.12%)	0.5419 (+0.88%)	0.4308 (+0.62%)	0.4148 (+1.39%)
	LlaMA-3.2-1B	0.9423 (−2.02%)	0.9842 (+3.29%)	0.8808 (+7.09%)	0.7349 (+15.9%)	0.9418 (+0.25%)	0.8968 (+4.90%)	0.4843 (+3.42%)	0.2871 (+1.06%)	0.3330 (+1.26%)	0.5038 (−2.93%)	0.4826 (+5.80%)	0.4182 (+1.72%)
	LlaMA-3.2-3B	0.9798 (+1.73%)	0.9734 (+2.21%)	0.8629 (+5.30%)	0.6417 (+6.58%)	0.9418 (+0.25%)	0.8799 (+3.21%)	0.4505 (+0.04%)	0.2719 (−0.46%)	0.3169 (−0.35%)	0.5083 (−2.48%)	0.4826 (+5.80%)	0.4060 (+0.51%)
	LlaMA-3.3-70B	0.9514 (−1.11%)	0.9707 (+1.94%)	0.7760 (−3.39%)	0.5676 (−0.83%)	0.8020 (−13.7%)	0.8135 (−3.42%)	0.5398 (+8.97%)	0.2738 (−0.27%)	0.3270 (+0.66%)	0.5397 (+0.66%)	0.5801 (+15.5%)	0.4521 (+5.11%)
Developments													
Fine-tuning	LlaMA-2-7B-chat-hf	0.9659	0.9636	0.8550	0.8860	0.9536	0.9248	0.2688	0.2043	0.3181	0.4155	0.3612	0.3136
	Vicuna-1.5-7B	0.9968 (+3.09%)	0.9937 (+3.01%)	0.7628 (−9.22%)	0.7992 (−8.68%)	0.9140 (−3.96%)	0.8933 (−3.15%)	0.3321 (+6.33%)	0.1489 (−5.54%)	0.3581 (+4.00%)	0.4209 (+0.54%)	0.4604 (+9.92%)	0.3441 (+3.05%)
	Wizardmath-7B	0.8487 (−11.7%)	0.9383 (−2.53%)	0.6720 (−18.3%)	0.7846 (−10.1%)	0.8257 (−12.7%)	0.8139 (−11.1%)	0.4049 (+13.6%)	0.2360 (+3.17%)	0.4409 (+12.2%)	0.4607 (+4.52%)	0.3993 (+3.81%)	0.3884 (+7.48%)
	Llama-2-7B-hf	0.6918 (−27.4%)	0.8354 (−12.8%)	0.5242 (−33.0%)	0.8093 (−7.67%)	0.7749 (−17.8%)	0.7271 (−19.7%)	0.4626 (+19.3%)	0.3657 (+16.1%)	0.4919 (+17.3%)	0.3611 (−5.44%)	0.4684 (+10.7%)	0.4299 (+11.6%)
Pruning	Sheared-LlaMA1.3B	0.6329 (−33.3%)	0.8232 (−14.0%)	0.5013 (−35.3%)	0.7989 (−8.71%)	0.6832 (−27.0%)	0.6879 (−23.6%)	0.6036 (+33.4%)	0.4374 (+23.3%)	0.5483 (+23.0%)	0.5286 (+11.3%)	0.5848 (+22.3%)	0.5405 (+22.7%)
	Sheared-LlaMA1.3B-pruned	0.2612 (−70.4%)	0.6781 (−28.5%)	0.3235 (−53.1%)	0.6188 (−26.7%)	0.4304 (−52.3%)	0.4624 (−46.2%)	0.8336 (+56.4%)	0.5806 (+37.6%)	0.6668 (+34.8%)	0.6672 (+25.1%)	0.6402 (+27.9%)	0.6777 (+36.4%)
	Sheared-LlaMA2.7B-pruned	0.6664 (−29.9%)	0.8491 (−11.4%)	0.5012 (−35.3%)	0.7650 (−12.1%)	0.6758 (−27.7%)	0.6915 (−23.3%)	0.5406 (+27.1%)	0.3740 (+16.9%)	0.5650 (+24.6%)	0.4914 (+7.59%)	0.5100 (+14.8%)	0.4962 (+18.2%)
	Sheared-LlaMA2.7B	0.4036 (−56.2%)	0.7400 (−22.3%)	0.3438 (−51.1%)	0.6612 (−22.4%)	0.5103 (−44.3%)	0.5318 (−39.3%)	0.7138 (+44.5%)	0.4800 (+27.5%)	0.6750 (+35.6%)	0.6056 (+19.0%)	0.5984 (+23.7%)	0.6146 (+30.1%)

Table 11: The detection performance (measured in AUROC) of LogRank and Entropy on EvoBench.

LLMs	Version Time/ Version Name	DetectGPT						LRR					
		Xsum	Writing	PubMed	SocialMedia	PeerRead	Avg.	Xsum	Writing	PubMed	SocialMedia	PeerRead	Avg.
Updates													
GPT-4o	2024-05-13	0.5182	0.8413	0.4455	0.4921	0.4643	0.5523	0.7032	0.7057	0.6719	0.4871	0.6213	0.6378
	2024-08-06	0.4785 (−3.97%)	0.7437 (−9.76%)	0.4558 (+1.03%)	0.5058 (+1.37%)	0.4450 (−1.93%)	0.5258 (−2.65%)	0.6655 (−3.77%)	0.6443 (−6.14%)	0.6601 (−1.18%)	0.5065 (+1.94%)	0.5744 (−4.69%)	0.6102 (−2.77%)
	2024-11-20	0.2796 (−23.8%)	0.7283 (−11.3%)	0.4636 (+1.81%)	0.5469 (+5.48%)	0.4668 (+0.25%)	0.4970 (−5.52%)	0.6201 (−8.31%)	0.6486 (−5.71%)	0.6593 (−1.26%)	0.5000 (+1.29%)	0.5696 (−5.17%)	0.5995 (−3.83%)
	Latest	0.2796 (−23.8%)	0.6685 (−17.2%)	0.4704 (+2.49%)	0.5108 (+1.87%)	0.4539 (−1.04%)	0.4766 (−7.56%)	0.6139 (−8.93%)	0.6611 (−4.46%)	0.6480 (−2.39%)	0.5103 (+2.32%)	0.5879 (−3.34%)	0.6042 (−3.36%)
GPT-4o-mini	2024-07-18	0.4990 (−1.92%)	0.7712 (−7.01%)	0.4474 (+0.19%)	0.5008 (+0.87%)	0.4902 (+2.59%)	0.5417 (−1.06%)	0.7239 (+2.07%)	0.6918 (−1.39%)	0.6569 (−1.50%)	0.5236 (+3.65%)	0.5818 (−3.95%)	0.6356 (−0.22%)
GPT-4	2023-06-13	0.3582	0.6876	0.3563	0.3688	0.3566	0.4255	0.7280	0.6961	0.6811	0.4311	0.5677	0.6208
	2023-11-06	0.5470 (+18.8%)	0.7465 (+5.89%)	0.4392 (+8.29%)	0.4385 (+6.97%)	0.4951 (+13.8%)	0.5333 (+10.7%)	0.7020 (−2.60%)	0.5723 (−12.3%)	0.6838 (+0.27%)	0.4747 (+4.36%)	0.6275 (+5.98%)	0.6121 (−0.87%)
	2024-01-25	0.5225 (+16.4%)	0.7256 (+3.80%)	0.4187 (+6.24%)	0.4772 (+10.8%)	0.4818 (+12.5%)	0.5252 (+9.97%)	0.6974 (−3.06%)	0.7293 (+3.32%)	0.6573 (−2.38%)	0.4677 (+3.66%)	0.6338 (+6.61%)	0.6371 (+1.63%)
	2024-04-09	0.5240 (+16.5%)	0.7389 (+5.13%)	0.4031 (+4.68%)	0.4856 (+11.6%)	0.4526 (+9.60%)	0.5208 (+9.53%)	0.6788 (−4.92%)	0.6120 (−8.41%)	0.6684 (−1.27%)	0.4918 (+6.07%)	0.6037 (+3.60%)	0.6109 (−0.99%)
Claude-Sonnet	2024-02-29	0.6532	0.8303	0.4666	0.3497	0.6006	0.5801	0.8454	0.8629	0.7312	0.3942	0.7993	0.7266
	2024-06-20	0.6551 (+0.19%)	0.8853 (+5.50%)	0.4738 (+0.72%)	0.4700 (+12.0%)	0.5718 (−2.88%)	0.6112 (+3.11%)	0.8236 (−2.18%)	0.8143 (−4.86%)	0.6860 (−4.52%)	0.4499 (+5.57%)	0.7341 (−6.52%)	0.7016 (−2.50%)
	2024-10-22	0.5885 (−6.47%)	0.8112 (−1.91%)	0.5158 (+4.92%)	0.4866 (+13.6%)	0.5704 (−3.02%)	0.5945 (+1.44%)	0.7560 (−8.94%)	0.6939 (−16.9%)	0.6557 (−7.55%)	0.4221 (+2.79%)	0.6230 (−17.6%)	0.6301 (−9.65%)
Claude-Haiku	2024-03-07	0.5728	0.7993	0.4760	0.4571	0.5511	0.5713	0.9439	0.9340	0.7639	0.5228	0.8174	0.7964
	2024-10-22	0.2622 (−31.0%)	0.4994 (−29.9%)	0.4924 (+1.64%)	0.3608 (−9.63%)	0.0210 (−53.0%)	0.3272 (−24.4%)	0.6454 (−29.8%)	0.5772 (−35.6%)	0.6065 (−15.7%)	0.3321 (−19.0%)	0.4288 (−38.8%)	0.5180 (−27.8%)
Claude-Opus	2024-02-29	0.5400 (−3.28%)	0.7742 (−2.51%)	0.5045 (+2.85%)	0.4575 (+0.04%)	0.5582 (+0.71%)	0.5669 (−0.44%)	0.8799 (−6.40%)	0.8756 (−5.84%)	0.7322 (−3.17%)	0.5164 (−0.64%)	0.7735 (−4.39%)	0.7555 (−4.09%)
Qwen	Qwen1.5-7B	0.4416	0.7407	0.3880	0.5055	0.6606	0.5473	0.9200	0.9508	0.6103	0.5260	0.7808	0.7576
	Qwen2-7B	0.6027 (+16.1%)	0.7513 (+1.06%)	0.4386 (+5.06%)	0.5067 (+0.12%)	0.6234 (−3.72%)	0.5845 (+3.73%)	0.9866 (+6.66%)	0.9850 (+3.42%)	0.6436 (+3.33%)	0.4996 (−2.64%)	0.7939 (+1.31%)	0.7817 (+2.42%)
	Qwen2.5-7B	0.5188 (+7.72%)	0.6895 (−5.12%)	0.4255 (+3.75%)	0.5146 (+0.91%)	0.6406 (−2.00%)	0.5578 (+1.05%)	0.9826 (+6.26%)	0.9819 (+3.11%)	0.6673 (+5.70%)	0.4999 (−2.61%)	0.7989 (+1.81%)	0.7861 (+2.85%)
LlaMA3	LlaMA-3.1-8B	0.5006	0.7926	0.4803	0.5480	0.6603	0.5964	0.9435	0.9057	0.6842	0.5301	0.9031	0.7933
	LlaMA-3.1-70B	0.5933 (+9.27%)	0.8360 (+4.34%)	0.4429 (−3.74%)	0.4987 (−4.93%)	0.6765 (+1.62%)	0.6095 (+1.31%)	0.9313 (−1.22%)	0.9416 (+3.59%)	0.6669 (−1.73%)	0.4895 (−4.06%)	0.8832 (−1.99%)	0.7825 (−1.08%)
	LlaMA-3.2-1B	0.4201 (−8.05%)	0.7794 (−1.32%)	0.4736 (−0.67%)	0.6325 (+8.45%)	0.6274 (−3.29%)	0.5866 (−0.98%)	0.9443 (+0.08%)	0.9827 (+7.70%)	0.7754 (+9.12%)	0.6932 (+16.3%)	0.9235 (+2.04%)	0.8638 (+7.05%)
	LlaMA-3.2-3B	0.5017 (+0.11%)	0.8217 (+2.91%)	0.4506 (−2.97%)	0.5150 (−3.30%)	0.6274 (−3.29%)	0.5833 (−1.31%)	0.9725 (+2.90%)	0.9612 (+5.55%)	0.7428 (+5.86%)	0.5544 (+2.43%)	0.9235 (+2.04%)	0.8309 (+3.76%)
	LlaMA-3.3-70B	0.5910 (+9.04%)	0.8195 (+2.69%)	0.4478 (−3.25%)	0.5073 (−4.07%)	0.5968 (−6.35%)	0.5925 (−0.39%)	0.9510 (+0.75%)	0.9366 (+3.09%)	0.7000 (+1.58%)	0.5253 (−0.48%)	0.7636 (−13.9%)	0.7753 (−1.80%)
Developments													
Fine-tuning	LlaMA-2-7B-chat-hf	0.4919	0.7052	0.5017	0.7118	0.6650	0.6151	0.9579	0.9436	0.7791	0.8559	0.9300	0.8933
	Vicuna-1.5-7B	0.4796 (−1.23%)	0.6116 (−9.36%)	0.3794 (−12.2%)	0.6358 (−7.60%)	0.5986 (−6.64%)	0.5410 (−7.41%)	0.9863 (+2.84%)	0.9815 (+3.79%)	0.7012 (−7.79%)	0.7583 (−9.76%)	0.8787 (−5.13%)	0.8612 (−3.21%)
	Wizardmath-7B	0.5770 (+8.51%)	0.6076 (−9.76%)	0.4606 (−4.11%)	0.6652 (−4.66%)	0.5814 (−8.36%)	0.5784 (−3.68%)	0.8240 (−13.3%)	0.9186 (−2.50%)	0.6076 (−17.1%)	0.8084 (−4.75%)	0.7438 (−18.6%)	0.7805 (−11.2%)
	Llama-2-7B-hf	0.4816 (−1.03%)	0.4501 (−25.5%)	0.5197 (+1.80%)	0.6717 (−4.01%)	0.5523 (−11.2%)	0.5351 (−8.00%)	0.7628 (−19.5%)	0.8694 (−7.42%)	0.5288 (−25.0%)	0.8418 (−1.41%)	0.7444 (−18.5%)	0.7494 (−14.3%)
Pruning	Sheared-LlaMA1.3B	0.4112 (−8.07%)	0.3640 (−34.1%)	0.4510 (−5.07%)	0.5996 (−11.2%)	0.3168 (−34.8%)	0.4285 (−18.6%)	0.7537 (−20.4%)	0.8866 (−5.70%)	0.5291 (−25.0%)	0.8554 (−0.05%)	0.7000 (−23.0%)	0.7450 (−14.8%)
	Sheared-LlaMA1.3B-pruned	0.0297 (−46.2%)	0.0538 (−65.1%)	0.2284 (−27.3%)	0.3150 (−39.6%)	0.1406 (−52.4%)	0.1535 (−46.1%)	0.5272 (−43.0%)	0.7888 (−15.4%)	0.4092 (−36.9%)	0.7304 (−12.5%)	0.5318 (−39.8%)	0.5975 (−29.5%)
	Sheared-LlaMA2.7B-pruned	0.5573 (+6.54%)	0.6327 (−7.25%)	0.4669 (−3.48%)	0.6359 (−7.59%)	0.4899 (−17.5%)	0.5565 (−5.86%)	0.9317 (−2.62%)	0.9638 (+2.02%)	0.5611 (−21.8%)	0.7985 (−5.74%)	0.6615 (−26.8%)	0.7833 (−11.0%)
	Sheared-LlaMA2.7B	0.4398 (−5.21%)	0.4220 (−28.3%)	0.4388 (−6.29%)	0.4282 (−28.3%)	0.2548 (−41.0%)	0.3967 (−21.8%)	0.7510 (−20.6%)	0.8904 (−5.32%)	0.5631 (−21.6%)	0.7493 (−10.6%)	0.5490 (−38.1%)	0.7006 (−19.2%)

Table 12: The detection performance (measured in AUROC) of DetectGPT and LRR on EvoBench.

LLMs	Version Time/ Version Name	NPR						DNA-GPT					
		Xsum	Writing	PubMed	SocialMedia	PeerRead	Avg.	Xsum	Writing	PubMed	SocialMedia	PeerRead	Avg.
Updates													
GPT-4o	2024-05-13	0.8971	0.9688	0.7346	0.5756	0.8256	0.8003	0.9945	0.8281	0.8276	0.5976	0.6781	0.7852
	2024-08-06	0.8696 (−2.75%)	0.9447 (−2.41%)	0.6992 (−3.54%)	0.5932 (+1.76%)	0.7667 (−5.89%)	0.7747 (−2.57%)	0.9934 (−0.11%)	0.8608 (+3.27%)	0.8223 (−0.53%)	0.6188 (+2.12%)	0.6856 (+0.75%)	0.7962 (+1.10%)
	2024-11-20	0.7202 (−17.6%)	0.8998 (−6.90%)	0.6969 (−3.77%)	0.5747 (−0.09%)	0.7831 (−4.25%)	0.7349 (−6.54%)	0.9858 (−0.87%)	0.7175 (−11.0%)	0.7868 (−4.08%)	0.6178 (+2.02%)	0.6634 (−1.47%)	0.7543 (−3.09%)
	Latest	0.6963 (−20.0%)	0.9075 (−6.13%)	0.7010 (−3.36%)	0.5820 (+0.64%)	0.8242 (−0.14%)	0.7422 (−5.81%)	0.9866 (−0.79%)	0.7532 (−7.49%)	0.7993 (−2.83%)	0.6422 (+4.46%)	0.6728 (−0.53%)	0.7708 (−1.44%)
GPT-4o-mini	2024-07-18	0.9062 (+0.91%)	0.9700 (+0.12%)	0.7304 (−0.42%)	0.6058 (+3.02%)	0.8042 (−2.14%)	0.8033 (+0.30%)	0.9978 (+0.33%)	0.8663 (+3.82%)	0.8208 (−0.68%)	0.6645 (+6.69%)	0.6385 (−3.96%)	0.7976 (+1.24%)
GPT-4	2023-06-13	0.9262	0.9776	0.7611	0.4937	0.7511	0.7819	0.9944	0.8725	0.8918	0.6562	0.7490	0.8328
	2023-11-06	0.9177 (−0.85%)	0.9300 (−4.76%)	0.7317 (−2.94%)	0.5480 (+5.43%)	0.8282 (+7.71%)	0.7911 (+0.92%)	0.9943 (−0.01%)	0.7876 (−8.49%)	0.8037 (−8.81%)	0.6586 (+0.24%)	0.6259 (−12.3%)	0.7740 (−5.88%)
	2024-01-25	0.8940 (−3.22%)	0.9435 (−3.41%)	0.7036 (−5.75%)	0.5360 (+4.23%)	0.8245 (+7.34%)	0.7803 (−0.16%)	0.9921 (−0.23%)	0.8002 (−7.23%)	0.8409 (−5.09%)	0.6516 (−0.46%)	0.6622 (−8.68%)	0.7894 (−4.34%)
	2024-04-09	0.8332 (−9.30%)	0.9120 (−6.56%)	0.7172 (−4.39%)	0.5764 (+8.27%)	0.7919 (+4.08%)	0.7661 (−1.58%)	0.9908 (−0.36%)	0.7880 (−8.45%)	0.8211 (−7.07%)	0.5970 (−5.92%)	0.6008 (−14.8%)	0.7595 (−7.32%)
Claude-Sonnet	2024-02-29	0.9459	0.9804	0.8150	0.5836	0.9320	0.8514	0.9700	0.8582	0.8403	0.7355	0.6106	0.8029
	2024-06-20	0.9692 (+2.33%)	0.9884 (+0.80%)	0.7718 (−4.32%)	0.6605 (+7.69%)	0.9030 (−2.90%)	0.8586 (+0.72%)	0.9861 (+1.61%)	0.8061 (−5.21%)	0.8159 (−2.44%)	0.7250 (−1.05%)	0.5549 (−5.57%)	0.7776 (−2.53%)
	2024-10-22	0.9016 (−4.43%)	0.9265 (−5.39%)	0.6760 (−13.9%)	0.5056 (−7.80%)	0.7789 (−15.3%)	0.7577 (−9.37%)	0.9548 (−1.52%)	0.7238 (−13.4%)	0.7418 (−9.85%)	0.6681 (−6.74%)	0.5468 (−6.38%)	0.7271 (−7.59%)
Claude-Haiku	2024-03-07	0.9952	0.9993	0.8614	0.7486	0.9389	0.9087	0.9972	0.9328	0.8388	0.7669	0.6745	0.8420
	2024-10-22	0.8396 (−15.5%)	0.9237 (−7.56%)	0.6439 (−21.7%)	0.3892 (−35.9%)	0.4641 (−47.4%)	0.6521 (−25.6%)	0.9965 (−0.07%)	0.8578 (−7.50%)	0.7293 (−10.9%)	0.7031 (−6.38%)	0.8217 (+14.7%)	0.8217 (−2.04%)
Claude-Opus	2024-02-29	0.9703 (−2.49%)	0.9604 (−3.89%)	0.8234 (−3.80%)	0.7159 (−3.27%)	0.8943 (−4.46%)	0.8729 (−3.58%)	0.9939 (−0.33%)	0.8727 (−6.01%)	0.8165 (−2.23%)	0.7698 (+0.29%)	0.6188 (−5.57%)	0.8143 (−2.77%)
Qwen	Qwen1.5-7B	0.9243	0.9945	0.6641	0.4897	0.9076	0.7960	0.9269	0.9034	0.7222	0.5074	0.4761	0.7072
	Qwen2-7B	0.9929 (+6.86%)	0.9974 (+0.29%)	0.7376 (+7.35%)	0.4622 (−2.75%)	0.8939 (−1.37%)	0.8168 (+2.08%)	0.8856 (−4.13%)	0.8692 (−3.42%)	0.8168 (+9.46%)	0.5209 (+1.35%)	0.5375 (+6.14%)	0.7260 (+1.88%)
	Qwen2.5-7B	0.9894 (+6.51%)	0.9955 (+0.10%)	0.7878 (+12.3%)	0.4686 (−2.11%)	0.8998 (−0.78%)	0.8282 (+3.22%)	0.9009 (−2.60%)	0.9100 (+0.66%)	0.7907 (+6.85%)	0.5205 (+1.31%)	0.5256 (+4.95%)	0.7295 (+2.23%)
LlaMA3	LlaMA-3.1-8B	0.9926	0.9839	0.8476	0.6866	0.9680	0.8957	0.9987	0.8939	0.7784	0.7232	0.6304	0.8049
	LlaMA-3.1-70B	0.9957 (+0.31%)	0.9989 (+1.50%)	0.8104 (−3.72%)	0.6693 (−1.73%)	0.9584 (−0.96%)	0.8865 (−0.92%)	0.9968 (−0.19%)	0.8807 (−1.32%)	0.7627 (−1.57%)	0.7361 (+1.29%)	0.6182 (−1.22%)	0.7989 (−0.60%)
	LlaMA-3.2-1B	0.9580 (−3.46%)	0.9948 (+1.09%)	0.9205 (+7.29%)	0.8617 (+17.5%)	0.9650 (−0.30%)	0.9400 (+4.43%)	0.9636 (−3.51%)	0.9378 (+4.39%)	0.8158 (+3.74%)	0.7512 (+2.80%)	0.7056 (+7.52%)	0.8348 (+2.99%)
	LlaMA-3.2-3B	0.9913 (−0.13%)	0.9870 (+0.31%)	0.9132 (+6.56%)	0.7698 (+8.32%)	0.9650 (−0.30%)	0.9253 (+2.95%)	0.9943 (−0.44%)	0.9210 (+2.71%)	0.8154 (+3.70%)	0.7655 (+4.23%)	0.6584 (+2.80%)	0.8309 (+2.60%)
	LlaMA-3.3-70B	0.9958 (+0.32%)	1.0000 (+1.61%)	0.7923 (−5.53%)	0.6881 (+0.15%)	0.9042 (−6.38%)	0.8761 (−1.97%)	0.9955 (−0.32%)	0.8865 (−0.74%)	0.7878 (+0.94%)	0.7308 (+0.76%)	0.6533 (+2.29%)	0.8108 (+0.59%)
Developments													
Fine-tuning	LlaMA-2-7B-chat-hf	0.9722	0.9904	0.9055	0.8963	0.9693	0.9467	0.6808	0.6330	0.4911	0.6396	0.6380	0.6165
	Vicuna-1.5-7B	0.9958 (+2.36%)	0.9956 (+0.52%)	0.8008 (−10.4%)	0.8805 (−1.58%)	0.9439 (−2.54%)	0.9233 (−2.34%)	0.9745 (+29.3%)	0.9305 (+29.7%)	0.8064 (+31.5%)	0.6490 (+0.94%)	0.6355 (−0.25%)	0.7992 (+18.2%)
	Wizardmath-7B	0.9205 (−5.17%)	0.9804 (−1.00%)	0.7354 (−17.0%)	0.8490 (−4.73%)	0.8412 (−12.8%)	0.8653 (−8.14%)	0.7476 (+6.68%)	0.7582 (+12.5%)	0.6969 (+20.5%)	0.6857 (+4.61%)	0.5447 (−9.33%)	0.6866 (+7.01%)
	Llama-2-7B-hf	0.5632 (−40.9%)	0.6202 (−37.0%)	0.4642 (−44.1%)	0.7395 (−15.6%)	0.6582 (−31.1%)	0.6091 (−33.7%)	0.6343 (−4.65%)	0.7310 (+9.80%)	0.4957 (+0.46%)	0.7474 (+10.7%)	0.7711 (+13.3%)	0.6759 (+5.94%)
Pruning	Sheared-LlaMA1.3B	0.6548 (−31.7%)	0.8888 (−10.1%)	0.4846 (−42.0%)	0.8634 (−3.29%)	0.7072 (−26.2%)	0.7198 (−22.7%)	0.7706 (+8.98%)	0.6626 (+2.96%)	0.5856 (+9.45%)	0.7255 (+8.59%)	0.5360 (−10.2%)	0.6561 (+3.96%)
	Sheared-LlaMA1.3B-pruned	0.4118 (−56.0%)	0.8077 (−18.2%)	0.3524 (−55.3%)	0.7530 (−14.3%)	0.4593 (−51.0%)	0.5568 (−38.9%)	0.9213 (+24.0%)	0.7903 (+15.7%)	0.6991 (+20.8%)	0.8669 (+22.7%)	0.6539 (+1.59%)	0.7863 (+16.9%)
	Sheared-LlaMA2.7B-pruned	0.4480 (−52.4%)	0.8672 (−12.3%)	0.3664 (−53.9%)	0.7543 (−14.2%)	0.5196 (−44.9%)	0.5911 (−35.5%)	0.8383 (+15.7%)	0.7618 (+12.8%)	0.5870 (+9.59%)	0.7831 (+14.3%)	0.5455 (−9.25%)	0.7031 (+8.66%)
	Sheared-LlaMA2.7B	0.7049 (−26.7%)	0.9154 (−7.50%)	0.5047 (−40.0%)	0.8191 (−7.72%)	0.6520 (−31.7%)	0.7192 (−22.7%)	0.7496 (+6.88%)	0.6811 (+4.81%)	0.5794 (+8.83%)	0.6743 (+3.47%)	0.4655 (−17.2%)	0.6300 (+1.35%)

Table 13: The detection performance (measured in AUROC) of NPR and DNA-GPT on EvoBench.