BEYOND RAW DETECTION SCORES: MARKOV-INFORMED CALIBRATION FOR BOOSTING MACHINE-GENERATED TEXT DETECTION

Anonymous authorsPaper under double-blind review

ABSTRACT

While machine-generated texts (MGTs) offer great convenience, they also pose risks such as disinformation and phishing, highlighting the need for reliable detection. Metric-based methods, which extract statistically distinguishable features of MGTs, are often more practical than complex model-based methods that are prone to overfitting. Given their diverse designs, we first place representative metric-based methods within a unified framework, enabling a clear assessment of their advantages and limitations. Our analysis identifies a core challenge across these methods: the token-level detection score is easily biased by the inherent randomness of the MGTs generation process. To address this, we theoretically and empirically reveal two relationships of context detection scores that may aid calibration: Neighbor Similarity and Initial Instability. We then propose a Markov-informed score calibration strategy that models these relationships using Markov random fields, and implements it as a lightweight component via a mean-field approximation, allowing our method to be seamlessly integrated into existing detectors. Extensive experiments in various real-world scenarios, such as cross-LLM and paraphrasing attacks, demonstrate significant gains over baselines with negligible computational overhead. The code is available at https://anonymous.4open.science/r/MRF-Enhance.

1 Introduction

In recent years, generative AI, represented by large language models (LLMs) (Achiam et al., 2023; Radford et al., 2019), has advanced rapidly, and the machine-generated texts (MGTs) they produce often match human writing in fluency, coherence, and diversity. While this technological breakthrough offers immense opportunities, it has also triggered widespread societal concerns, such as the spread of disinformation (Vykopal et al., 2024), the violation of intellectual property rights (Yu et al., 2023b), and phishing attacks (Hong, 2012). Therefore, the research and development of MGT detection technologies hold significant theoretical and practical value in uncovering the distinct patterns of generated text and ensuring a trustworthy AI environment.

An effective detection method is to identify LLM's watermarks (Hou et al., 2024), but this requires injecting watermarks into the LLM, which is often impractical due to high access permissions. Therefore, passive detection methods, including model-based and metric-based methods, have garnered significant attention. Model-based methods, which use a set of human- and machine-generated texts to train a binary classifier, such as OpenAI detector (Solaiman et al., 2019), ChatGPT detector (Guo et al., 2023), SeqXGPT (Wang et al., 2023), and CoCo (Liu et al., 2022). However, such models are often too complex, leading to overfitting to the training data. Instead, metric-based methods exploit the inherent statistical biases of LLM to discriminate MGTs, which is model-agnostic and has better generalization properties. These methods use metrics such as log-likelihood, log-rank, and entropy. Furthermore, methods such as DetectGPT (Mitchell et al., 2023), DNA-GPT (Yang et al., 2024), and SimLLM (Nguyen-Son et al., 2024) detect MGTs by comparing the differences between a given text and a perturbed, regenerated, or continued text from an alternative model.

Obviously, metric-based methods exhibit diverse designs. Therefore, this paper first systematically examines several representative approaches, including Log-Likelihood (Solaiman et al., 2019), Log-

Rank (Mitchell et al., 2023), Entropy (Gehrmann et al., 2019), DetectGPT (Mitchell et al., 2023), Fast-DetectGPT (Bao et al., 2024), and DNA-GPT (Yang et al., 2024), and situates them within a unified framework (**Section 2**). Our analysis reveals that they share a threshold-based detection criterion, with only minor differences, such as the inclusion of auxiliary data (e.g., perturbed texts). This offers a theoretical basis for understanding their core mechanisms, strengths, and limitations.

Based on the unified framework, we summarize the core challenge of metric-based methods (**Section 2**): the imprecision of token-level detection scores. Specifically, since these methods make decisions based on a threshold, their effectiveness is directly tied to score precision. However, randomness introduced by the LLM sampling mechanism can violate their underlying assumptions, leading to biased scores with low discrimination, as shown in "w/o refine" in Fig. 1 (more results can be found in Appendix E.1). Moreover, they tend to derive an overall detection score by naively aggregating token-level scores, failing to correct the underlying imprecision. Therefore, calibrating the token-level detection score is essential for improving overall detection performance.

Given that detection scores are tied to tokens and the LLM generation process induces dependencies among tokens (Achiam et al., 2023), context tokens' detection scores may have relationships that are easy to overlook. Revealing and modeling these relationships may help calibrate these scores. Accordingly, we theoretically and empirically reveal two relationships among detection scores of context tokens (Section 3): Neighbor Similarity, where adjacent tokens exhibit similar detection scores

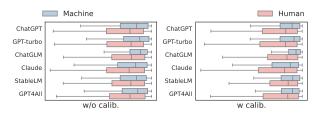


Figure 1: Distribution of token scores obtained by the DetectGPT method without and with score calibration in the Essay dataset. The proposed calibration method enhances the discriminative nature of the token scores.

kens exhibit similar detection scores, and **Initial Instability**, where the detection scores of initial tokens are unstable.

Finally, building on these two relationships, we propose a Markov-informed score calibration method to enhance MGT detection (**Section 4**). Our method models the identified relationships through Markov random fields and, via a mean-field approximation, implements it as a lightweight iterative neural network. As shown by the more discriminative detection scores in the "w refine" method in Fig. 1, the proposed method boosts the discriminative nature of the scores. Notably, our method can be seamlessly stacked on top of existing detectors without architectural changes, providing flexibility. Compared with complex model-based approaches, our method introduces only a negligible 2×2 parameterization, making its computational delay negligible and less prone to overfitting. Extensive experiments demonstrate our method's enhanced effectiveness. Our contributions can be summarized as follows:

- We view existing metric-based detection methods through a unified lens, which facilitates precise comparison and enables potential improvements.
- We theoretically and empirically demonstrate that token-level detection scores exhibit neighbor similarity and initial instability, offering avenues for improved detection.
- We propose a Markov-informed score calibration method and, via a mean-field approximation, implement it as a lightweight component that can be seamlessly integrated into existing detectors to further unlock their potential.
- We conduct extensive experiments across three datasets to demonstrate superior performance in diverse scenarios, including cross-LLM, cross-domain, mixed-text, and paraphrase attacks.

2 A Unified Perspective on Metric-Based Detection

Although model-based methods have shown competitive potential in specific domains, they are often too complex, leading to a tendency to overfit their training data. This limitation requires them to retrain or fine-tune for newly released LLMs, which hinders their generalizability. In contrast, metric-based methods extract discriminative features from MGT, and their model-agnostic nature provides superior generalization potential. Given the diverse implementations of representative metric-based

Table 1: Comparing existing metric-based methods from a unified view. Here, s is the text to be detected containing N tokens, s' is the perturbed text generated by DetectGPT, \tilde{s} and \hat{s} are the regenerated texts of Fast-DetectGPT and DNA-GPT, respectively. Function $\mu(\cdot)$ and $\sigma(\cdot)$ represent the mean and standard deviation of the given set, respectively.

Method	Perturb	Sample	Score	Detection
Log-likelihood	/	s	$\frac{1}{N-1} \sum_{t=2}^{N} \log p(s_t s_{< t})$	$ score > \epsilon $
Log-Rank	/	s	$\frac{1}{N-1} \sum_{t=2}^{N} \text{rank}(p(s_t s_{< t}))$	$score > \epsilon$
Entropy	/	s	$-\frac{1}{N-1} \sum_{t=2}^{N} \sum_{v \in V} p(v s_{< t}) \log p(v s_{< t})$	$score > \epsilon$
DetectGPT	$\{s_1', s_2',, s_n'\}$	s	$\frac{\frac{1}{N-1}\sum_{t=2}^{N}p(s_{t} s_{< t}) - \mu\left(\left\{\frac{1}{N-1}\sum_{t=2}^{N}p(s_{i,t}' s_{i,< t}')\right\}_{i}\right)}{\sigma\left(\left\{\frac{1}{N-1}\sum_{t=2}^{N}p(s_{i,t}' s_{i,< t}')\right\}_{i}\right)}$	$score > \epsilon$
Fast-DetectGPT	/	$\{s, \tilde{s}_1,, \tilde{s}_n\}$	$\frac{\frac{1}{N-1} \sum_{t=2}^{N} p(\hat{s}_{t} s_{< t}) - \mu\left(\left\{\frac{1}{N-1} \sum_{t=2}^{N} p(\hat{s}_{i,t} s, \tilde{s}_{i,< t})\right\}_{i}\right)}{\sigma\left(\left\{\frac{1}{N-1} \sum_{t=2}^{N} p(\tilde{s}_{i,t} s, \tilde{s}_{i,< t})\right\}_{i}\right)}$	$score > \epsilon$
DNA-GPT	/	$\{s, \hat{s}_1,, \hat{s}_n\}$	$\frac{1}{N-1} \sum_{t=2}^{N} \log p(s_t s_{1:t-1}) - \frac{1}{n} \sum_{i=1}^{n} \frac{1}{N_t - 1} \sum_{t=2}^{N_t} \log p(\hat{s}_{i,t} \hat{s}_{i,1:t-1})$	$score > \epsilon$

methods such as Log-Likelihood (Solaiman et al., 2019), Log-Rank (Mitchell et al., 2023), Entropy (Gehrmann et al., 2019), DetectGPT (Mitchell et al., 2023), Fast-DetectGPT (Bao et al., 2024), and DNA-GPT (Yang et al., 2024), we first provide a systematic examination of them from a unified perspective. This facilitates a deeper understanding of their mechanisms and allows for a fair comparison of their strengths and weaknesses. As illustrated in Table 1, we compare these methods across data, score calculation, and detection dimensions.

- Data. Log-likelihood, Log-Rank, and Entropy are computationally efficient as they rely solely on the original input text s. However, the detection error based on a single text may be large, because the randomness inherent in the LLM sampling mechanism may cause the MGT to deviate from these methods' underlying assumptions, e.g., Log-Rank assumes that the generated tokens have high rankings. In contrast, DetectGPT, Fast-DetectGPT, and DNA-GPT incorporate multiple perturbed (i.e., s') or regenerated (i.e., s and s) samples, which mitigates the errors caused by randomness. However, this comes at the cost of a significant increase in computational overhead.
- Score Calculation. Although these methods appear to calculate scores differently, they all tend to directly aggregate token scores to obtain the final text score, typically through summation. As discussed, the randomness introduced by the LLM generation process may cause token-level scores to be biased. Therefore, the direct aggregation of these potentially imprecise token scores may result in an inaccurate final detection score.
- **Detection**. These methods all employ threshold-based detection mechanisms, whose effectiveness relies heavily on the accuracy of their calculated scores. As previously discussed, two factors compromise this accuracy: (1) the inherent randomness of LLM-generated text introduces bias into score calculation, and (2) direct score aggregation fails to mitigate this bias. As a result, their detection performance is often unsatisfactory.

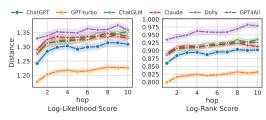
In summary, to enhance detection, existing methods incorporate more textual information (e.g., regenerated texts in DetectGPT and Fast-DetectGPT) and use different score calculation strategies (e.g., the Likelihood difference between the detected and regenerated text in DNAGPT). However, as we discussed, they fail to address the underlying token-level errors caused by inherent randomness, limiting their detection potential. Considering that detection scores are tied to tokens and the LLMs' generative mechanism induces dependencies among tokens, revealing and modeling the relationships between tokens may help correct score errors and thus improve detection effectiveness.

3 RELATION BETWEEN CONTEXTUAL TOKEN-LEVEL DETECTION SCORES

To understand the relationship between context tokens' detection scores, following existing work (Liu et al., 2023), we consider the token generation process of a simplified model: a single-layer transformer model with single-head attention:

$$x_{t+1} = \mathcal{F}(a_t)$$
, where $a_t = \operatorname{softmax} \left(1/t \cdot x_t W_Q W_K^\top X_{t-1}^\top \right) X_{t-1} W_V W_O$. (1)

 $x_t \in \mathbb{R}^{1 \times d}$ is the embedding of token s_t , and d denotes the embedding dimension. The matrix $X_{t-1} \in \mathbb{R}^{(t-1) \times d}$ is stacked by the embeddings x_1, \ldots, x_{t-1} , where the j-th row is x_j .



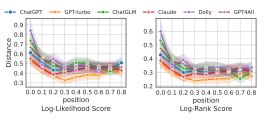


Figure 2: The detection score distances of neighbors at different hops in the Essay dataset. Log-likelihood and Log-Rank score are used here.

Figure 3: The score distances of 1-hop neighbors at different positions in the Essay dataset. Log-likelihood and Log-Rank score are used here.

 $W_Q, W_K, W_V \in \mathbb{R}^{d \times d}$ and $W_O \in \mathbb{R}^{d \times d}$ are the attention weights. Following the attention block, an MLP block, denoted as $\mathcal{F}: \mathbb{R}^{1 \times d} \to \mathbb{R}^{1 \times d}$, is applied, and it is a two-layer network with skip connections:

$$\mathcal{F}(x) = x + W_2 \operatorname{relu}(W_1 x).$$

As shown in Table 1, the detection score of token s_t is usually the function of s_t (i.e., x_i), and x_i is related to the attention scores $\alpha_{t-1} = \operatorname{softmax} \left(1/t \cdot x_t W_Q W_K^\top X_{t-1}^\top \right)$ in Eq. (1). The following theorem will reveal the relationship between attention scores, which in turn help us understand the relationship between detection scores of context tokens.

Theorem 1. Let $\lambda_K, \lambda_Q, \lambda_V, \lambda_O$ be the largest singular values of parameters W_K, W_Q, W_V, W_O , respectively, and let $W = W_V W_O W_Q W_K^{\top}$. For the transformer defined in Eq. (1), assuming normalized inputs $(\|x_t\|_2 = 1 \text{ for all } t)$ and constants $c, \epsilon > 0$, consider $a_t x_{t+1}^{\top} \geq (1 - \delta) \|a_t\|_2$ with $\delta \leq \left(\frac{c\epsilon}{\lambda_Q \lambda_K \lambda_V \lambda_O}\right)^2$. If x_ℓ satisfies $x_\ell W x_\ell^{\top} \geq c$ and $x_\ell W x_\ell \geq \epsilon^{-1} \max_{j \in [\ell], j \neq \ell} x_j W x_\ell^{\top}$, then

$$\alpha_{t+1,l} \leq \frac{\exp\left(C_l \cdot \alpha_{t,l} + \eta\right)}{\exp\left(C_l \cdot \alpha_{t,l} + \eta\right) + \sum_{j \neq l} \exp\left(C_j \cdot \alpha_{t,j} - \eta\right)},$$

$$\alpha_{t+1,l} \geq \frac{\exp\left(C_l \cdot \alpha_{t,l} - \eta\right)}{\exp\left(C_l \cdot \alpha_{t,l} - \eta\right) + \sum_{j \neq l} \exp\left(C_j \cdot \alpha_{t,j} + \eta\right)},$$

where

$$C_j = rac{x_j W x_j^T}{t \left| a_t \right|_2}, and \ \eta = rac{(1+\sqrt{2})\epsilon x_j W x_j^T}{(t+1) \left| a_t \right|_2}.$$

The proof can be found in Appendix C. This theorem establishes the upper and lower bounds for the attention score at step t+1, which are determined by the attention scores at step t. This finding creates a positive feedback loop analogous to simulated annealing (Kirkpatrick et al., 1983), where high attention scores at the current step lead to high scores at the next, and vice versa. Meanwhile, it ensures that the attention score remains stable over time throughout the generation process.

As the detection scores of tokens are the function of their attention values, this theorem reflects two relationships between detection scores of context tokens: (1) **Neighbor Similarity**. Compared with other tokens, detection scores of neighbor tokens tend to be more similar. (2) **Initial Instability**. Compared with the subsequent tokens, the detection scores of initial tokens are unstable (i.e., fluctuate greatly). In addition to theoretical results, we further verify these two relationships empirically.

First, to empirically validate the neighbor similarity property, we evaluated the detection score difference between the token s_t and its k-hop neighbors (i.e., tokens s_{t-k} and s_{t+k} if existing). As illustrated in Fig. 2 (more results can be found in Appendix E.2), there is a clear positive correlation between the detection score distance and the hop, and adjacent tokens have the highest detection score similarity; thereby providing empirical evidence for our theoretical finding.

Second, to validate the initial instability property, we analyzed the detection score difference between adjacent tokens at different percentage positions. Fig. 3 illustrates that the score difference is substantially larger for tokens at the beginning of the text and progressively decreases, eventually stabilizing. More results can be found in Appendix E.2. Considering our established finding on neighbor similarity (i.e., adjacent scores should be highly similar), a high detection score difference at the sequence beginning indicates a significant instability in detection scores of initial tokens.

4 MARKOV-INFORMED DETECTION SCORE CALIBRATION

Based on the revealed relationships, this section uses an MRF to capture them (Section 4.1) and adopts the mean field approximation to model the MRF model as a lightweight component stacked on existing detectors to calibrate detection scores (Section 4.2), thereby enhancing detection.

4.1 Markov Random Field for MGT Detection

We capture these two types of relationships by modeling the joint probability distribution of text's token detection scores through pairwise Markov Random Fields (pMRF). Specifically, for each token s_t in text s, we assign a binary random variable y_{s_t} , where $y_{s_t}=0$ and $y_{s_t}=1$ indicate a human- or machine-generated token 1 , respectively, as measured by the detection score of the token. Let y_s denote the label set for all tokens in text s, the pMRF over these tokens can be formalized as a Gibbs distribution: $P(y_s)=\frac{1}{Z}\exp(-E(s,y_s))$, where Z is a normalizing constant and $E(s,y_s)$ is the energy function. Our objective is to maximize the posterior probability of the token labels y_s by minimizing the global energy function $E(s,y_s)$. The energy function typically consists of two components: the unary potential Ψ_U and the pairwise potential Ψ_P :

$$E(s, y_s) = \sum_{t=1}^{N} \Psi_U(s_t, y_{s_t}) + \sum_{t=1}^{N} \sum_{s_j \in \mathcal{N}(s_t)} \Psi_P(y_{s_t}, y_{s_j}),$$

where $\mathcal{N}(s_t)$ denotes the adjacent tokens of token s_t , i.e., $\mathcal{N}(s_t) = \{s_{t-1}, s_{t+1}\}$ (if existing).

Unary potential $\Psi_U(s_t, y_{s_t})$ quantifies the cost of assigning label y_{s_t} to token s_t . We let $\Psi_U(s_t, y_{s_t}) = -\log p(s_t)$, where $p(s_t)$ is the output probability from the original detector. For detectors without probability output, it is measured by the 0-1 normalized score of token s_t .

Pairwise potential $\Psi_P\left(y_{s_t},y_{s_j}\right)$ models the similarity in detection scores between adjacent tokens. A penalty is applied if two adjacent tokens are assigned different labels; otherwise, a reward is given. This enforces label smoothness and captures the neighbor similarity property:

$$\Psi_P(y_{s_t}, y_{s_j}) = w \cdot (2 \cdot I(y_{s_t} \neq y_{s_j}) - 1), \tag{2}$$

where $I(\cdot)$ is the indicator function, and the reward and penalty factor $w \ge 0$. This implies an energy penalty of w when adjacent tokens have different labels; otherwise, the reward is -w.

To model the initial instability property, we introduce a positional weighting function $\beta(t)$ in the binary potential. This function assigns lower weights to binary potentials at earlier positions, thereby mitigating the amplification of energy errors caused by unstable initial neighbor tokens. In this paper, we define the positional weighting function $\beta(t)$ as a Sigmoid function to ensure a smooth transition of weights, and the revised binary potential is then given by:

$$\Psi_P(y_{s_i}, y_{s_j}) = \beta(j) \cdot w \cdot (2 \cdot I(y_{s_i} \neq y_{s_j}) - 1), \text{ with } \beta(t) = \frac{1}{1 + exp(-(t - t_0))},$$
 (3)

where t_0 is the weighted center, effectively suppressing the pairwise potential of tokens before t_0 .

4.2 MEAN FIELD APPROXIMATE IN MGT DETECTION

Given the MRF model, this subsection details how to model it as a lightweight component stacked on the original detector through mean field approximation theory, thereby enhancing detection.

In the MRF posterior probability $P(y_s) = \frac{1}{Z} \exp(-E(s,y_s))$, the partition function $Z = \sum_{y_s} exp(-E(s,y_s))$, which is obtained by adding over all possible combinations of y_s . For a text with N tokens, there are 2^N combinations, making exact computation of $P(y_s|s)$ infeasible. Inspired by existing work (Deng et al., 2022), we employ mean-field theory for approximate inference. Its core idea is to use a simpler, factorized distribution $Q(y_s) = \prod_{t=1}^N Q_{s_t}(y_{s_t})$ to approximate the true joint distribution $P(y_s)$, achieved by minimizing the KL divergence between these two distributions:

$$D(Q||P) = \mathbb{E}_{y_s \sim Q} \left[\log Q \left(y_s \right) \right] - \mathbb{E}_{y_s \sim Q} \left[\log P \left(y_s \right) \right]$$

$$= \sum_{t=1}^{N} \mathbb{E}_{y_{s_t} \sim Q_{s_t}} \left[\log Q_{s_t} \left(y_{s_t} \right) \right] + \mathbb{E}_{y_s \sim Q} \left[E \left(s, y_s \right) \right] + \log Z.$$

$$(4)$$

¹Note that token labels are not absolute but depend on the context in which they appear. For example, "the" can be a human token or a machine token depending on the text.

 The first equation is the definition of KL divergence, and the second is obtained by substituting $P(y_s)$ and $Q(y_s)$. Then, we define a Lagrangian composed of all terms involving $Q_{s_t}(y_{s_t})$ in D(Q||P):

$$L_{s_{t}}(Q) = Q_{s_{t}}(y_{s_{t}}) \log Q_{s_{t}}(y_{s_{t}}) + \mathbb{E}_{y_{s} \sim Q} \left[E\left(s, y_{s}\right) \right] + \lambda \left(\sum_{y_{s_{t}}} Q_{s_{t}}(y_{s_{t}}) - 1 \right).$$
 (5)

Here, the term involving Lagrange multiplier λ assures that Q_{s_i} is a proper probability distribution. Now we take derivatives of Eq. (5) with respect to $Q_{s_t}(y_{s_t})$ and set the corresponding derivative to 0, then we get the optimal $Q_{s_t}(y_{s_t})$:

$$Q_{s_t}(y_{s_t}) = \frac{1}{z} exp(\Psi_U(s_t, y_{s_t}) - \sum_{s_i \in \mathcal{N}(s_t)} \mathbb{E}_{y_{s_i} \sim Q_{s_i}} \left[\Psi_P(y_{s_t}, y_{s_i}) \right]). \tag{6}$$

We can express the single token calculation in Eq. (6) as a matrix form of multiple token calculations, which involves three steps:

- Unary potential calculation. Corresponding to the unary potential $\Psi_U(s_t, y_{s_t})$ of token s_t being $-\log p(s_i)$, we can express the unary potential of text s in matrix form $-\log H$, where the i-th row of H corresponds to token s_i 's prior probability and the two columns correspond to identity labels (HGT and MGT), that is, H = [1 p(s), p(s)], where $p(s) = [p(s_1), ..., p(s_N)]^{\top}$.
- Pairwise potential calculation. For the revised pairwise potential (i.e., Eq. (3)), we can define the weighted adjacency matrix as A^{corr} , where $A^{corr}_{t,t+1} = \beta(t+1)$ for all $t=0,...,N-2, A^{corr}_{t-1,t} = \beta(t-1)$ for all t=1,...,N-1, and 0 otherwise. Then the matrix form of the weighted pairwise potential is $A^{corr}Q\begin{bmatrix} -w & w \\ w & -w \end{bmatrix}$. This update strategy enforces that all relationships receive the same reward or penalty. However, the influence of MGT neighbors and HGT neighbors is intuitively different, so this kind of reward and punishment mechanism with the same weight will limit the expressive ability of MRF. To this end, we relax the weights for different relationships and get the pairwise potential as $A^{corr}Q(W_{mrf}\odot\begin{bmatrix} -1 & 1 \\ 1 & -1 \end{bmatrix}$), where $W_{mrf}\in\mathbb{R}^{2\times 2}_+$.
- Normalization. The operator $\frac{1}{z}exp(\cdot)$ in Eq. (6) can be naturally modeled as Softmax function.

In summary, we get the following update rule for tokens' detection scores:

$$Q = \operatorname{softmax} \left(\log H - A^{corr} Q \left(W_{mrf} \odot \begin{bmatrix} -1 & 1 \\ 1 & -1 \end{bmatrix} \right) \right). \tag{7}$$

Notably, Eq. (7) shows that the computation of Q relies on Q itself; hence, iterative computation is required. Initializing the initial Q to H, we can get the iterative version as follows:

$$Q^{t} = \operatorname{softmax} \left(\log Q^{t-1} - A^{corr} Q^{t-1} \left(W_{mrf} \odot \begin{bmatrix} -1 & 1 \\ 1 & -1 \end{bmatrix} \right) \right), \text{ where } Q^{0} = H.$$
 (8)

In addition to using position weight function $\beta(t)$ in pairwise potential to reduce the impact of initial unstable scores, we also use this position weight function on the final calibrated scores Q^T of T iterations to reduce their impact on detection:

$$Q_{final} = [\beta(1), ..., \beta(N)] \odot Q^{T}.$$
 (9)

The MRF-informed calibration component is then directly stacked on the original detector to calibrate the token detection scores. Specifically, if the de-

Algorithm 1 Markov-informed Enhancement Framework

- 1: **Input:** Text s to be detected, the original detector $f_1 \circ f_2$, iteration steps T, MRF weights W.
- 2: Construct A^{corr} based on s.
- 3: Get each token s_i 's detection score from the detection score calculation module f_1 , and set $Q^0 = H$.
- 4: **for** t = 0 **to** T 1 **do**
- 5: Update Q according to Eq. (8).
- 6: end for
- 7: Calculate Q_{final} according to Eq. (9).
- 8: **Return** detection score $f_2(Q_{final})$ of text s.

tector is formalized as a combination of the detection score calculation module f_1 and the detection module f_2 , i.e., $f(s) = f_1 \circ f_2(s)$, the MRF-informed component of Eq. (9) is defined as f_{mrf} , then the enhanced detector is $f_{enh.}(s) = f_1 \circ f_{mrf} \circ f_2(s)$. The complete inference process is shown in Alg. 1. To learn weights W, we use supervised training: $\mathcal{L} = -\sum_{s \in \mathcal{D}_{train}} (Y_s \log f_{enh.}(s) + (1 - Y_s) \log (1 - f_{enh.}(s)))$.

325 326

327

328

330

331

332

333

334

335336337

338

340 341

342

343

344

345346347

348

349

350

351 352

353 354

355

356

357

358

359

360

361

362

364

365

366

367

368

369

370

371

372

373 374

375

376

377

Figure 4: The performance improvement of the proposed method on Likelihood and Log-Rank. Values greater than 0 indicate an enhanced effect.

writing xsum Test Domain writing xsum Test Domain

Log-Rank

writing xsum Test Domain

Likelihood

Computational Complexity. The MRF layer can be computed using sparse-dense matrix multiplication. For a text containing N tokens, the number of iterations is T, resulting in a computational complexity of $\mathcal{O}(NT)$, which is negligible for calculating the detection score. We will also empirically verify the efficiency of our method in Appendix E.8.

5 EXPERIMENTS

writing xsum Test Domain

Dataset. We conducted experiments on three widely used public datasets: Essay (Verma et al., 2024), Reuters (Verma et al., 2024), and DetectRL (Wu et al., 2024). The Essay and Reuters datasets collect machine text generated by GPT4All, ChatGPT, ChatGPT-turbo, ChatGLM, Dolly, and Claude. The DetectRL dataset not only includes pure machine-generated text by Google-PaLM, ChatGPT, and Llama-2-70b, but its unique mixed, paraphrased, and cross-domain texts also allow us to more comprehensively evaluate the model's performance in complex real-world scenarios. For a complete description of the datasets, please refer to Appendix D.1.

Baselines. We select the following metric-based methods for comparison and enhancement: Log-Likelihood (Likelihood) (Solaiman et al., 2019), Log-Rank (Mitchell et al., 2023), Entropy (Gehrmann et al., 2019), DetectGPT (Mitchell et al., 2023), Fast-DetectGPT (FastGPT) (Bao et al., 2024), and DNA-GPT (Yang et al., 2024). The versions equipped with the proposed method are defined with 'M' suffix, e.g., Likelihood-M and Log-Rank-M. Furthermore, although we focus on metric-based methods, we also compare with model-based methods ChatGPT-D (Guo et al., 2023) and MPU (Tian et al., 2024). Their details can be found in Appendix D.2.

Metrics. First, as a binary classification problem, we use the area under the receiver operating characteristic curve (AUROC). Second, following (Tufts et al., 2024; Fraser et al., 2025; Hans et al., 2024), we recognize the negative impact of misclassifying human text as machine-generated text. Therefore, another important evaluation metric is the true positive rate (TPR) at a low false positive rate (FPR). Specifically, we measure the TPR at an FPR of 1%, denoting this as TPR@FPR-1%.

5.1 Performance Comparison

In this section, we evaluate the enhancement effectiveness of the proposed method in various real-world scenarios, including cross-LLM, cross-domain, detecting mixed machine text, and resisting paraphrase attacks. Details of these scenarios are provided in Appendix D.3.

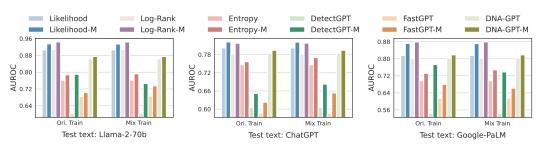


Figure 5: Detection performance concerning AUROC under different LLM mixed texts. All detectors are trained on Llama-2-70b texts.



Figure 6: Detection performance concerning AUROC under different paraphrasing texts (Dipper and Polish). All detectors are trained on Llama-2-70b texts.

Performance across Different LLMs. Table 2 reports AUROC for detectors trained on GPT4All (Essay) and Llama-2-70b (DetectRL) and evaluated on texts from various LLMs. TPR@FPR-1% results are provided in Table 3 in the Appendix. Besides, results of detectors trained on other LLM texts and on the Reuters dataset appear in Tables 4–15 of the Appendix. The proposed method yields significant gains for nearly all baselines. For example, on Essay, it raises Likelihood from 52.4% to 77.86% (+25.46%). Encouragingly, our method also benefits weak detectors: DetectGPT significantly improves from 0.15% to 37.18% (+37.03%), suggesting that the assumptions underlying its scoring are reasonable and that underperformance mainly stemmed from score estimation error. Finally, detection performance on texts from the same LLM is not always superior to cross-LLM. For example, on Essay, performance on ChatGLM is particularly strong and even exceeds the intra-LLM case of GPT4All, possibly because ChatGLM texts are more discriminative.

Performance across Different Domains. We evaluate cross-domain performance in four high-risk domains: arXiv, Writing Prompts, XSum, and Yelp Reviews. Results are summarized in Fig. 4; additional enhanced results for more detectors under cross-domain settings are provided in Fig. 19 and Fig. 20 of the Appendix. In most settings, detectors equipped with our strategy show substantially stronger cross-domain generalization. We attribute this to calibrating the detection score with only a few carefully designed parameters (a 2x2 reward–penalty coefficient matrix), which helps prevent overfitting and thereby improves out-of-domain detection.

Performance against Mixed Texts. In practice, human–AI collaboration is pervasive, leading to widespread mixed human–machine text. We therefore evaluate the proposed strategy for mixed-text detection. Two training strategies are considered: training on the original text and training on the mixed text. Results concerning AUROC are shown in Fig. 5, with TPR@FPR-1% shown in Fig. 21 of the Appendix. In most settings, the proposed strategy enhances the detector's ability to recognize mixed-text. Moreover, comparing training on original versus mixed text shows that simply training on mixed text does not improve the detection ability of mixed text, highlighting the focus of mixed text detection research.

Performance against Paraphrasing Attacks. Prior work (Sadasivan et al., 2023) shows that MGT detection is typically vulnerable to paraphrasing attacks, where an adversary rewrites a passage without altering its semantics to evade detection. We therefore assess the robustness gains of our framework on the Dipper and Polish paraphrase attacks provided by DetectRL. The results in Fig. 6 and Fig. 22 clearly indicate that, even in adversarial settings, our strategy yields encouraging improvements in robustness against paraphrasing attacks.

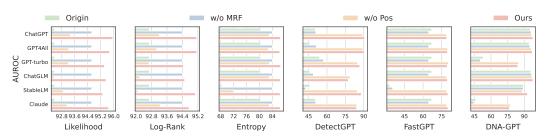


Figure 7: Ablation results on the Essay dataset. The y-axis represents the LLM text on which the detector was trained, and the x-axis represents the average performance across LLMs.

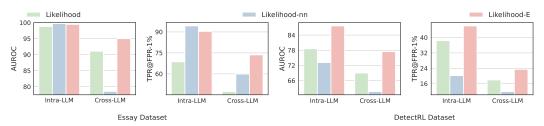


Figure 8: Comparison with NN-based calibration methods. The detector used is Likelihood.

5.2 ABLATION STUDY

We introduce an MRF layer and a positional weighting function to model neighbor similarity and initial instability, respectively. In this section, we verify their effectiveness through ablation studies, denoted as "w/o MRF" and "w/o Pos". Results on the Essay dataset are shown in Fig. 7, with additional results in Fig. 23 and Fig. 24 of the Appendix. We can find that removing either component significantly drops performance, while retaining only one still outperforms the baseline detector, underscoring their designs' rationality. Besides, the contribution of each component varies by detector type. For single-text detectors like Likelihood and Log-Rank, the positional weighting function provides the most improvement by addressing instability in initial scores. However, for methods that aggregate multiple text scores, such as DetectGPT and FastGPT, this instability is already partially mitigated, making the MRF-based score calibration the primary source of gains.

5.3 MRF vs. Neural Network

This paper proposes a Markov-informed calibration method to model the relationship of detection scores of context tokens. To highlight the rationale of this design, we compare it with methods that directly use neural networks to calibrate scores. A simple three-layer neural network is used here, defined with the "nn" suffix. The comparison results on Likelihood are shown in Fig. 8, and more results can be found in Fig. 29 and Fig. 30 of the Appendix. Although NN-based methods exhibit competitive performance in some settings for intra-LLMs, their generalization ability on cross-LLMs drops significantly. This suggests that it does not truly learn the general ability to correct scores, but rather overfits the training data. In contrast, our strategy shows good generalization.

6 Conclusion

This paper has systematically examined representative metric-based detectors within a unified framework, revealing a core challenge: the inherent randomness of the LLM generation process leads to inaccurate detection scores, and naive aggregation of existing methods fails to fix this. Therefore, we have theoretically and empirically established two key properties of these scores: neighbor similarity and initial instability. Building on these insights, we have proposed a Markovinformed score calibration method that captures token relationships and corrects the biased scores produced by base detectors. Extensive experiments show substantial and consistent performance gains of the proposed method. Admittedly, our strategy is only applicable to metric-based methods that utilize token-level features. We will try to improve model-based methods in future work.

ETHICS STATEMENT

This paper proposes a Markov-informed calibration strategy to enhance machine-generated text detection and mitigate the potential risks posed by machine-generated text, including disinformation and phishing. Our work does not involve ethical issues such as dataset releases, potentially harmful insights, potential conflicts of interest and sponsorship, discrimination/bias/fairness concerns, privacy and security issues, legal compliance, and research integrity issues.

REPRODUCIBILITY STATEMENT

Our code is available at https://anonymous.4open.science/r/MRF-Enhance, and all datasets used in our experiments (Essay, Reuters, and DetectRL) are publicly available for download. In addition, we provide detailed implementation details in the Appendix, including data partitioning, the fixed seeds, the learning rate, the training batch size, and the two parameters t_0 and T introduced by the proposed strategy, to ensure reproducibility of our work.

REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. arXiv preprint arXiv:2303.08774, 2023.
- 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 *ICLR*, 2024.
- Leyan Deng, Chenwang Wu, Defu Lian, Yongji Wu, and Enhong Chen. Markov-driven graph convolutional networks for social spammer detection. *IEEE Transactions on Knowledge and Data Engineering*, 35(12):12310–12322, 2022.
- Kathleen C Fraser, Hillary Dawkins, and Svetlana Kiritchenko. Detecting ai-generated text: Factors influencing detectability with current methods. *Journal of Artificial Intelligence Research*, 82: 2233–2278, 2025.
- Sebastian Gehrmann, Hendrik Strobelt, and Alexander M Rush. Gltr: Statistical detection and visualization of generated text. arXiv preprint arXiv:1906.04043, 2019.
- GPTZero. Gptzero official website. [Online], 2023. https://gptzero.me.
- Chenchen Gu, Xiang Lisa Li, Percy Liang, and Tatsunori Hashimoto. On the learnability of water-marks for language models. *arXiv preprint arXiv:2312.04469*, 2023.
- Biyang Guo, Xin Zhang, Ziyuan Wang, Minqi Jiang, Jinran Nie, Yuxuan Ding, Jianwei Yue, and Yupeng Wu. How close is chatgpt to human experts? comparison corpus, evaluation, and detection. *arXiv preprint arXiv:2301.07597*, 2023.
- Abhimanyu Hans, Avi Schwarzschild, Valeriia Cherepanova, Hamid Kazemi, Aniruddha Saha, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Spotting llms with binoculars: Zeroshot detection of machine-generated text. In *International Conference on Machine Learning*, pp. 17519–17537. PMLR, 2024.
- Jason Hong. The state of phishing attacks. *Communications of the ACM*, 55(1):74–81, 2012.
- Abe Hou, Jingyu Zhang, Tianxing He, Yichen Wang, Yung-Sung Chuang, Hongwei Wang, Lingfeng Shen, Benjamin Van Durme, Daniel Khashabi, and Yulia Tsvetkov. Semstamp: A semantic watermark with paraphrastic robustness for text generation. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 4067–4082, 2024.
- Xiaomeng Hu, Pin-Yu Chen, and Tsung-Yi Ho. Radar: Robust ai-text detection via adversarial learning. *Advances in neural information processing systems*, 36:15077–15095, 2023.

- John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. A
 watermark for large language models. In *International Conference on Machine Learning*, pp. 17061–17084. PMLR, 2023.
- Scott Kirkpatrick, C Daniel Gelatt Jr, and Mario P Vecchi. Optimization by simulated annealing. *science*, 220(4598):671–680, 1983.
 - Kalpesh Krishna, Yixiao Song, Marzena Karpinska, John Wieting, and Mohit Iyyer. Paraphrasing evades detectors of ai-generated text, but retrieval is an effective defense. *Advances in Neural Information Processing Systems*, 36:27469–27500, 2023.
 - Linyang Li, Pengyu Wang, Ke Ren, Tianxiang Sun, and Xipeng Qiu. Origin tracing and detecting of llms. *arXiv preprint arXiv:2304.14072*, 2023.
 - Xiaoming Liu, Zhaohan Zhang, Yichen Wang, Hang Pu, Yu Lan, and Chao Shen. Coco: Coherence-enhanced machine-generated text detection under data limitation with contrastive learning. *arXiv* preprint arXiv:2212.10341, 2022.
 - Yepeng Liu and Yuheng Bu. Adaptive text watermark for large language models. In *Forty-first International Conference on Machine Learning*.
 - Zichang Liu, Aditya Desai, Fangshuo Liao, Weitao Wang, Victor Xie, Zhaozhuo Xu, Anastasios Kyrillidis, and Anshumali Shrivastava. Scissorhands: Exploiting the persistence of importance hypothesis for llm kv cache compression at test time. *Advances in Neural Information Processing Systems*, 36:52342–52364, 2023.
 - Shixuan Ma and Quan Wang. Zero-shot detection of llm-generated text using token cohesiveness. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 17538–17553, 2024.
 - Eric Mitchell, Yoonho Lee, Alexander Khazatsky, Christopher D Manning, and Chelsea Finn. Detectgpt: Zero-shot machine-generated text detection using probability curvature. In *International Conference on Machine Learning*, pp. 24950–24962. PMLR, 2023.
 - Hoang-Quoc Nguyen-Son, Minh-Son Dao, and Koji Zettsu. Simllm: Detecting sentences generated by large language models using similarity between the generation and its re-generation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 22340–22352, 2024.
 - Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
 - Vinu Sankar Sadasivan, Aounon Kumar, Sriram Balasubramanian, Wenxiao Wang, and Soheil Feizi. Can ai-generated text be reliably detected? *arXiv preprint arXiv:2303.11156*, 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. Release strategies and the social impacts of language models. *arXiv preprint arXiv:1908.09203*, 2019.
 - Jinyan Su, Terry Yue Zhuo, Di Wang, and Preslav Nakov. Detectllm: Leveraging log rank information for zero-shot detection of machine-generated text. In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
 - Yuchuan Tian, Hanting Chen, Xutao Wang, Zheyuan Bai, QINGHUA ZHANG, Ruifeng Li, Chao Xu, and Yunhe Wang. Multiscale positive-unlabeled detection of ai-generated texts. In *The Twelfth International Conference on Learning Representations*, 2024.
 - Brian Tufts, Xuandong Zhao, and Lei Li. An examination of ai-generated text detectors across multiple domains and models. In *Neurips Safe Generative AI Workshop 2024*, 2024.
 - Eduard Tulchinskii, Kristian Kuznetsov, Laida Kushnareva, Daniil Cherniavskii, Sergey Nikolenko, Evgeny Burnaev, Serguei Barannikov, and Irina Piontkovskaya. Intrinsic dimension estimation for robust detection of ai-generated texts. *Advances in Neural Information Processing Systems*, 36, 2024.

- Vivek Verma, Eve Fleisig, Nicholas Tomlin, and Dan Klein. Ghostbuster: Detecting text ghostwritten by large language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 1702–1717, 2024.
- Ivan Vykopal, Matúš Pikuliak, Ivan Srba, Robert Moro, Dominik Macko, and Mária Bieliková. Disinformation capabilities of large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 14830–14847, 2024.
- Pengyu Wang, Linyang Li, Ke Ren, Botian Jiang, Dong Zhang, and Xipeng Qiu. Seqxgpt: Sentence-level ai-generated text detection. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023.
- Junchao Wu, Runzhe Zhan, Derek Wong, Shu Yang, Xinyi Yang, Yulin Yuan, and Lidia Chao. Detectrl: Benchmarking llm-generated text detection in real-world scenarios. *Advances in Neural Information Processing Systems*, 37:100369–100401, 2024.
- Kangxi Wu, Liang Pang, Huawei Shen, Xueqi Cheng, and Tat-Seng Chua. Llmdet: A large language models detection tool. *CoRR*, 2023a.
- Yihan Wu, Zhengmian Hu, Hongyang Zhang, and Heng Huang. Dipmark: A stealthy, efficient and resilient watermark for large language models. 2023b.
- Xianjun Yang, Kexun Zhang, Haifeng Chen, Linda Petzold, William Yang Wang, and Wei Cheng. Zero-shot detection of machine-generated codes. *arXiv preprint arXiv:2310.05103*, 2023.
- 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*, 2024.
- KiYoon Yoo, Wonhyuk Ahn, and Nojun Kwak. Advancing beyond identification: Multi-bit water-mark for large language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 4031–4055, 2024.
- Xiao Yu, Yuang Qi, Kejiang Chen, Guoqiang Chen, Xi Yang, Pengyuan Zhu, Weiming Zhang, and Nenghai Yu. Gpt paternity test: Gpt generated text detection with gpt genetic inheritance. *CoRR*, 2023a.
- Xiao Yu, Kejiang Chen, Qi Yang, Weiming Zhang, and Nenghai Yu. Text fluoroscopy: Detecting llm-generated text through intrinsic features. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 15838–15846, 2024.
- Zhiyuan Yu, Yuhao Wu, Ning Zhang, Chenguang Wang, Yevgeniy Vorobeychik, and Chaowei Xiao. Codeipprompt: intellectual property infringement assessment of code language models. In *International conference on machine learning*, pp. 40373–40389. PMLR, 2023b.
- Haolan Zhan, Xuanli He, Qiongkai Xu, Yuxiang Wu, and Pontus Stenetorp. G3detector: General gpt-generated text detector. *arXiv preprint arXiv:2305.12680*, 2023.
- Ruisi Zhang, Shehzeen Samarah Hussain, Paarth Neekhara, and Farinaz Koushanfar. {REMARK-LLM}: A robust and efficient watermarking framework for generative large language models. In 33rd USENIX Security Symposium (USENIX Security 24), pp. 1813–1830, 2024.

The Use of Large Language Models

APPENDIX A More Discussion of the Proposed Method **B** Related Work Proof of Theorem 1 **D** Experimental Details **E** More Experimental Resuls More Results of Token Score Distribution before and after Enhancement E.3 E.4 E.5 More Results Comparing with Neural Network Calibration E.6 E.7 E.8

A MORE DISCUSSION OF THE PROPOSED METHOD

A.1 CONTRIBUTION

The contributions of this paper are multifaceted.

- We provide a unified perspective for understanding metric-based detection methods. The diversity of these methods (e.g., using metrics based on log-likelihood or log-rank, and introducing perturbed or regenerated text) makes comparison difficult. To this end, we re-examine these methods from three perspectives: data, score calculation, and detection. This analysis provides a precise definition for each method, facilitating comparison and potential improvements. Our analysis shows that these methods employ more reasonable evaluation metrics and incorporate additional contextual information to enhance detection. However, they all fail to address the underlying token-level errors caused by inherent randomness, which limits their detection potential. Therefore, our unified analysis encourages a more nuanced characterization of token-level detection scores, which will provide guidance for future improvements.
- We reveal the relationships between contextual token detection scores. While the generative mechanisms of LLMs introduce dependencies between tokens, these relationships remain unclear. To this end, we theoretically reveal two relationships between contextual token scores: neighboring similarity and initial instability, which are further validated through empirical experiments. Constraining these relationships during detection has the potential to mitigate the imprecision in score calculation caused by the inherent randomness of MGTs, which is crucial for the field of MGT detection.
- We propose a Markov-informed score calibration method to enhance MGT detection. This involves using Markov random fields to capture the revealed relationships and, through mean-field approximation, modeling the MRF model as a lightweight component that can be stacked on existing detectors to further unlock detection potential. It is worth noting that our main technical contribution and innovation lies not only in the specific implementation but also in the conceptual token-level score calibration. While the current implementation is based on Markov random fields, this is only one possible approach; alternatives include using sequence models or graph neural networks. This conceptual insight can inspire improving MGT detection.
- Extensive experiments consistently demonstrate the enhanced effectiveness of the proposed strategy. We empirically verify that it not only excels on a single task but also demonstrates strong capabilities in multiple complex and challenging real-world scenarios, including generalization across LLMs and domains, and robustness to mixed text and paraphrased text. Furthermore, the proposed enhancement component incurs negligible computational overhead compared to the original detector. This combination of effectiveness and efficiency provides a solid foundation for developing practically deployable enhancement solutions for AI-generated text detection.

A.2 LIMITATION

The core of the proposed method relies on modeling the relationships between context tokens to calibrate detection scores. Consequently, our method is not directly applicable to detection methods that do not provide this fine-grained, token-level output. For example, model-based methods typically generate a single document-level score, while some rely on features such as emotional expression or writing style. Future research could explore developing calibration models for different levels of feature granularity.

B RELATED WORK

Existing detection methods can be categorized into active watermark-based methods and passive model-based and metric-based methods.

B.1 WATERMARK-BASED DETECTION

Watermarking is a proactive defense technique that embeds verifiable information during the text generation stage, thereby enabling simple and reliable detection. RedList (Kirchenbauer et al., 2023)

is a model-agnostic watermarking method that dynamically partitions the vocabulary into a "greenlist" and "redlist" based on preceding context, slightly increasing the probability of sampling tokens from the greenlist. Subsequent works have made various improvements to this approach. For instance, SemStamp (Hou et al., 2024) introduces a sentence-level semantic hashing watermark to enhance robustness against paraphrasing attacks; DiPmark (Wu et al., 2023b) designs an unbiased watermark that does not alter the original output distribution. REMARK-LLM (Zhang et al., 2024) is a training-based watermarking method that employs a message encoding module to generate an encrypted token distribution for watermark embedding prior to inference. Beyond manually designed watermarks, directly leveraging language models to learn to generate watermarked text is also promising, including training student models (Gu et al., 2023) and semantically invariant watermarking models (Liu & Bu). In addition to the standard binary (0/1 bit) detection of AI-generated content, researchers have also explored multi-bit watermarks (Yoo et al., 2024) for embedding more information.

B.2 MODEL-BASED DETECTION

756

757

758

759

760

761

762

763

764

765

766

767

768 769

770 771

772

773

774

775

776

777

778

779

780

781

782

783

784 785

786 787

788 789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

Model-based methods represent a classical paradigm in detection, training a binary classifier on a dataset containing both human- and machine-generated texts. A series of works, such as OpenAI Detector (Solaiman et al., 2019), ChatGPT Detector (Guo et al., 2023), GPTZero (GPTZero, 2023), and G3 Detector (Zhan et al., 2023), collect texts generated by various LLMs to train a unified classifier. GPT-Pat (Yu et al., 2023a) finds that detectors trained solely on a single decoding strategy generalize poorly and enhances performance by utilizing mixed decoding strategies. In addition to original data, GLTR (Gehrmann et al., 2019) trains a simple logistic regression classifier by analyzing the predicted ranking of each word within its context. SeqXGPT (Wang et al., 2023) treats the sequence of logits as waveform signals for detection, while Sniffer (Li et al., 2023) uses the difference in logits from different models on the same text as features for detection and attribution. Beyond the data level, recent works have explored more advanced training strategies. For example, CoCo (Liu et al., 2022) introduces graph structures and contrastive learning; LLMDet (Wu et al., 2023a) leverages the perplexity of surrogate models as additional features; MPU (Tian et al., 2024) adopts a positive-unlabeled learning paradigm; and RADAR (Hu et al., 2023) incorporates adversarial training to enhance model robustness. The above methods generally assume a known text source, but when it is unknown, Ghostbuster (Verma et al., 2024) proposes training classifiers directly using texts generated from known surrogate models.

B.3 Metric-based Detection

Metric-based methods do not require training on specific datasets; instead, they directly leverage the inherent statistical biases or intrinsic properties of language model-generated text for distinction. The main advantage of such methods lies in their stronger generalization to new models and domains. Classic approaches in this category include the use of Log-Likelihood (Solaiman et al., 2019), Log-Rank (Mitchell et al., 2023), and Entropy (Gehrmann et al., 2019). DetectGPT (Mitchell et al., 2023) finds that AI-generated text typically lies in regions of negative curvature with respect to the model's log-probability function. By perturbing the text and observing changes in log-probability, it can effectively distinguish AI-generated text. Inspired by DetectGPT, Fast-DetectGPT (Bao et al., 2024) replaces log-probability with conditional probability curvature, significantly improving detection efficiency while maintaining performance. DetectLLM-LRR (Su et al.) proposes using the ratio of log-likelihood to log-rank for detection. Some works, such as DNA-GPT (Yang et al., 2024) and DetectGPT4Code (Yang et al., 2023), detect AI-generated text by comparing discrepancies between the original text and continuations generated by a surrogate model. PHD (Tulchinskii et al., 2024) observes that genuine human-written text possesses higher intrinsic dimensionality after encoder mapping. SimLLM (Nguyen-Son et al., 2024) is based on the observation that the similarity between the original text and its generated continuation is significantly higher than that between generated text and its re-generated version; thus, it estimates the similarity between an input sentence and its generated counterpart for detection. Given that existing methods struggle with out-of-distribution data, token coherence (Ma & Wang, 2024) has been shown to be a reliable metric since LLM-generated text usually exhibits higher token coherence than human-written text. Yu et al. (Yu et al., 2024) capture the intrinsic features of text by identifying layers with the greatest distributional differences when projecting into the vocabulary space, and using intrinsic rather than semantic features for detection has been demonstrated to yield better results.

C PROOF OF THEOREM 1

Theorem. I Let $\lambda_K, \lambda_Q, \lambda_V, \lambda_O$ be the largest singular values of parameters W_K, W_Q, W_V, W_O , respectively, and let $W = W_V W_O W_Q W_K^{\top}$. For the transformer defined in Eq. (1), assuming normalized inputs $(\|x_t\|_2 = 1 \text{ for all } t)$ and constants $c, \epsilon > 0$, consider $a_t x_{t+1}^{\top} \geq (1 - \delta) \|a_t\|_2$ with $\delta \leq \left(\frac{c\epsilon}{\lambda_Q \lambda_K \lambda_V \lambda_O}\right)^2$. If x_ℓ satisfies $x_\ell W x_\ell^{\top} \geq c$ and $x_\ell W x_\ell \geq \epsilon^{-1} \max_{j \in [\ell], j \neq \ell} x_j W x_\ell^{\top}$, then

$$\alpha_{t+1,l} \le \frac{\exp\left(C_l \cdot \alpha_{t,l} + \eta\right)}{\exp\left(C_l \cdot \alpha_{t,l} + \eta\right) + \sum_{j \ne l} \exp\left(C_j \cdot \alpha_{t,j} - \eta\right)},$$

$$\alpha_{t+1,l} \ge \frac{\exp\left(C_l \cdot \alpha_{t,l} - \eta\right)}{\exp\left(C_l \cdot \alpha_{t,l} - \eta\right) + \sum_{j \ne l} \exp\left(C_j \cdot \alpha_{t,j} + \eta\right)},$$

where

$$C_j = \frac{x_j W x_j^T}{t \left| a_t \right|_2}, \text{ and } \eta = \frac{(1 + \sqrt{2})\epsilon x_j W x_j^T}{(t+1) \left| a_t \right|_2}.$$

Proof. Our proof follows the proof of existing work (Liu et al., 2023). First, we introduce two necessary lemmas to help our proof.

Lemma 2 ((Liu et al., 2023)). Let $x_1, x_2 \in \mathbb{R}^{1 \times m}$ satisfies $||x_1||_2 = ||x_2||_2 = 1$ and $x_1 x_2^{\top} \ge 1 - \delta$ for some $\delta \in (0, 1)$. Then for all $y \in \mathbb{R}^{1 \times m}$ we have

$$|x_1 y^{\top} - x_2 y^{\top}| \le \sqrt{2\delta} ||y||_2$$

Lemma 3 ((Liu et al., 2023)). Let $\ell \in [t]$ be given. Suppose that $x_{\ell}Ax_{\ell}^{\top} > \epsilon^{-1} |x_{j}Ax_{\ell}^{\top}|$ for all $j \neq \ell$. Then we have

$$(\mathcal{S}(t)_{\ell} - \epsilon) x_{\ell}^{\top} a x_{\ell} \leq x_{\ell}^{\top} W_{K}^{\top} W_{Q} a_{t} \leq (\mathcal{S}(t)_{\ell} + \epsilon) x_{\ell}^{\top} a x_{\ell}$$

Based on these two lemmas, we can formally prove the theorem. Let $x_1 = \frac{a_t}{\|a_t\|}$, and $x_2 = x_{t+1}$. If $\frac{a_t x_{t+1}^\top}{\|a_t\|} \ge 1 - \delta$, using the conclusion of Lemma 2, we have

$$\left| \frac{a_t}{(t+1) \|a_t\|_2} W_Q W_K^T x_\ell^T - \frac{1}{t+1} x_{t+1} W_Q W_K^T x_\ell^T \right| \le \frac{\sqrt{2\delta}}{t+1} \left\| W_Q W_K^T x_\ell^T \right\|_2$$

Since λ_Q , λ_K are the maximum singular values, respectively. Then we have $\|W_Q W_K^\top x_\ell^\top\|_2 \le \lambda_Q \lambda_K \|x_\ell\|_2 = \lambda_Q \lambda_K$. This leads to:

$$\left| \frac{a_t}{(t+1) \|a_t\|_2} W_Q W_K^T x_\ell^T - \frac{1}{t+1} x_{t+1} W_Q W_K^T x_\ell^T \right| \le \frac{\sqrt{2\delta}}{t+1} \lambda_Q \lambda_K \tag{10}$$

Since

$$\|a_t\|_2 = \left\| \left(\sum_{j=1}^{t-1} \alpha_{t,j} x_j \right) W_V W_O \right\| \le \lambda_O \lambda_V \left\| \sum_{j=1}^{t-1} \alpha_{t,j} x_j \right\|_2 \le \lambda_O \lambda_V \sum_{j=1}^{t-1} \alpha_{t,j} \|x_j\|_2 = \lambda_O \lambda_V,$$

and the theorem assumes $\delta \leq \left(\frac{c\epsilon}{\lambda_Q \lambda_K \lambda_V \lambda_O}\right)^2$, substituting these into Eq. (10), we can obtain:

$$\frac{\sqrt{2\delta}}{t+1}\lambda_Q\lambda_K \le \frac{\sqrt{2}c\epsilon}{(t+1)\lambda_V\lambda_O} \le \frac{\sqrt{2}c\epsilon}{(t+1)\|a_t\|_2} \le \frac{\sqrt{2}\epsilon}{(t+1)\|a_t\|_2} x_\ell a x_\ell^T \tag{11}$$

The last inequality is obtained from $x_{\ell}ax_{\ell}^T \geq c$. Then combining Formula (10) and Formula (11), we have

$$\left| \frac{a_t}{(t+1) \|a_t\|_2} W_Q W_K^T x_\ell^T - \frac{1}{t+1} x_{t+1} W_Q W_K^T x_\ell^T \right| \le \frac{\sqrt{2}\epsilon}{(t+1) \|a_t\|_2} x_\ell a x_\ell^T$$

From Lemma 3, we have:

$$\left| \frac{a_{t} W_{Q} W_{K}^{\top} x_{\ell}^{\top}}{(t+1) \|a_{t}\|^{2}} - \frac{\alpha_{t,\ell} x_{\ell} W x_{\ell}^{\top}}{(t+1) \|a_{t}\|^{2}} \right| \leq \frac{\epsilon}{(t+1) \|a_{t}\|^{2}} x_{\ell}^{\top} a x_{\ell}$$

By the triangle inequality, we can combine the upper bounds:

$$\left| \frac{x_{t+1} W_{Q} W_{K}^{T} x_{\ell}^{T}}{t+1} - \frac{\alpha_{t,\ell} x_{\ell} W x_{\ell}^{T}}{(t+1) \|a_{t}\|_{2}} \right| \leq \left| \frac{x_{t+1} W_{Q} W_{K}^{T} x_{\ell}^{T}}{t+1} - \frac{a_{t} W_{Q} W_{K}^{T} x_{\ell}^{T}}{(t+1) \|a_{t}\|_{2}} \right| + \left| \frac{a_{t} W_{Q} W_{K}^{T} x_{\ell}^{T}}{(t+1) \|a_{t}\|_{2}} - \frac{\alpha_{t,\ell} x_{\ell} W x_{\ell}^{T}}{(t+1) \|a_{t}\|_{2}} \right|$$

$$\leq \frac{(1+\sqrt{2})\epsilon}{(t+1) \|a_{t}\|^{2}} x_{\ell}^{T} a x_{\ell}$$

$$(12)$$

Then, we rearrange the inequality, and we can obtain:

$$\frac{x_{\ell}Wx_{\ell}^{\top}}{(t+1)\left\|a_{t}\right\|_{2}}\left(\alpha_{t,\ell}-(1+\sqrt{2})\epsilon\right) \leq \frac{1}{t+1}x_{t+1}W_{Q}W_{K}^{\top}x_{\ell}^{\top} \leq \frac{x_{\ell}Wx_{\ell}^{\top}}{(t+1)\left\|a_{t}\right\|_{2}}\left(\alpha_{t,\ell}+(1+\sqrt{2}\epsilon)\right) \leq \frac{1}{t+1}x_{t+1}W_{Q}W_{K}^{\top}x_{\ell}^{\top} \leq \frac{x_{\ell}Wx_{\ell}^{\top}}{(t+1)\left\|a_{t}\right\|_{2}}\left(\alpha_{t,\ell}+(1+\sqrt{2}\epsilon)\right) \leq \frac{1}{t+1}x_{t+1}W_{Q}W_{K}^{\top}x_{\ell}^{\top} \leq \frac{x_{\ell}Wx_{\ell}^{\top}}{(t+1)\left\|a_{t}\right\|_{2}}\left(\alpha_{t,\ell}+(1+\sqrt{2}\epsilon)\right)$$

Now we give the lower and upper bounds of $\alpha_{t+1,\ell} = \operatorname{softmax}(1/(t+1) \cdot x_{t+1} W_Q W_K^\top X_t^\top)_l$. Upper bound. Let $\gamma_{t+1,\ell} = 1/(t+1) \cdot x_{t+1} W_Q W_K^\top x_\ell^\top$. For the softmax function $\alpha_{t+1,l} = \frac{\exp(\gamma_{t+1,\ell})}{\exp(\gamma_{t+1,\ell}) + \sum_{k \neq \ell} \exp(\gamma_{t+1,k})}$, to get the maximum value of $\alpha_{t+1,\ell}$, we need to (1) make the numerator as big as possible, which means $\gamma_{t+1,\ell}$ should take its maximum value $\frac{x_\ell W x_\ell^\top}{(t+1)\|a_t\|_2} \left(\alpha_{t,\ell} + (1+\sqrt{2})\epsilon\right)$, (2) make the denominator as small as possible, which means that all other values $\gamma_{t+1,k}$ (when $k \neq \ell$) should be minimized $\frac{x_k W x_k^\top}{(t+1)\|a_t\|_2} \left(\alpha_{t,k} - (1+\sqrt{2})\epsilon\right)$. Therefore,

$$\alpha_{t+1,l} \leq \frac{\frac{x_{\ell}Wx_{\ell}^{\top}}{(t+1)\|a_{t}\|_{2}} \left(\alpha_{t,\ell} + (1+\sqrt{2})\epsilon\right)}{\frac{x_{\ell}Wx_{\ell}^{\top}}{(t+1)\|a_{t}\|_{2}} \left(\alpha_{t,\ell} + (1+\sqrt{2})\epsilon\right) + \sum_{k \neq \ell} \frac{x_{k}Wx_{k}^{\top}}{(t+1)\|a_{t}\|_{2}} \left(\alpha_{t,k} - (1+\sqrt{2})\epsilon\right)}$$

Lower bound. Similarly, for the lower bound, we should (1) make the numerator as small as possible, which means $\gamma_{t+1,\ell}$ should take its minimum value $\frac{x_\ell W x_\ell^\top}{(t+1)\|a_t\|_2} \left(\alpha_{t,\ell} - (1+\sqrt{2})\epsilon\right)$, (2) make the denominator as large as possible, which means that all other values $\gamma_{t+1,k}$ (when $k \neq \ell$) should be maximized $\frac{x_k W x_k^\top}{(t+1)\|a_t\|_2} \left(\alpha_{t,k} + (1+\sqrt{2})\epsilon\right)$. Therefore,

$$\alpha_{t+1,l} \geq \frac{\frac{x_{\ell}Wx_{\ell}^{\top}}{(t+1)\|a_{t}\|_{2}}\left(\alpha_{t,\ell} - (1+\sqrt{2})\epsilon\right)}{\frac{x_{\ell}Wx_{\ell}^{\top}}{(t+1)\|a_{t}\|_{2}}\left(\alpha_{t,\ell} - (1+\sqrt{2})\epsilon\right) + \sum_{k \neq \ell} \frac{x_{k}Wx_{k}^{\top}}{(t+1)\|a_{t}\|_{2}}\left(\alpha_{t,k} + (1+\sqrt{2})\epsilon\right)}$$

The proof is completed.

D EXPERIMENTAL DETAILS

D.1 DATASETS

The details of the dataset used in the paper are as follows:

Essay (Verma et al., 2024). Each source of this dataset (human-written texts, various LLM-generated texts) contains 1,000 samples. The HGT portion comprises original IvyPanda essays that cover a wide array of subjects and academic levels, from high school through university. For the MGT portion, a tailored prompt was first crafted for each source essay using ChatGPT-turbo, and that prompt was then submitted to several LLMs, including GPT4All, ChatGPT, ChatGPT-turbo, ChatGLM, Dolly, and Claude, to generate machine-written essays. This workflow produced a diverse set of model-generated texts that remained aligned with the topics of their corresponding source documents.

- Reuters (Verma et al., 2024). Built on the Reuters 50–50 authorship benchmark, this dataset contains 1,000 articles from 50 writers, with each author contributing 20 pieces. Replicating the pipeline used for the essay corpus, the team first asked ChatGPT-turbo to invent a headline for every article. Those auto-generated headlines were then embedded into prompts and submitted to multiple LLMs, including ChatGPT, GPT-4, ChatGPT-turbo, ChatGLM, Dolly, and Claude, to create the machine-generated texts.
- **DetectRL** (Wu et al., 2024). In this dataset, the human-authored portion is drawn from four sources: arXiv abstracts dated 2002–2017, XSum news reports, Writing Prompts stories, and Yelp reviews. These types were chosen because they are especially vulnerable to producing convincing but misleading content when LLMs are misapplied. From each source, 2,800 human texts are selected as HGTs. The machine-generated texts are created using four widely used LLMs—GPT-3.5-turbo (ChatGPT), PaLM-2-bison (Google-PaLM), and Llama-2-70b. The dataset further models practical adversarial settings: (1) a paraphrasing attack that rewrites MGTs with the Dipper paraphraser (Krishna et al., 2023) and Polish paraphraser, and (2) a mixed-text condition where 1/4 of machine-generated sentences is randomly replaced with human-written content while the label remains "machine-generated."

D.2 BASELINES

A detailed description of the baselines used is shown below:

- **Likelihood** (Solaiman et al., 2019). It uses an LLM to calculate the log probability of each token in a text. The average of these probabilities gives a detection score. A higher score indicates a greater chance that the text was generated by LLMs.
- Log-Rank (Mitchell et al., 2023). Its detection score is created by first using an LLM to rank each token in a text based on its predicted order within a given context. The logarithm of each word's predicted rank is then calculated. The final score is an average of these values, and a lower score is a strong indicator of machine-generated text.
- Entropy (Gehrmann et al., 2019). Similar to Log-Rank, it calculates a score for a text by taking the average of each token's conditional entropy within its given context. A lower score suggests a higher likelihood that the text was generated by LLMs.
- **DetectGPT** (Mitchell et al., 2023). It determines if a text is machine-generated by measuring how small changes affect its log probability. The underlying idea is that text created by LLMs is already a high-probability output. So, when it is slightly altered, the new version is likely to have a lower log probability. In contrast, making similar small changes to human-written text does not consistently lower the log probability; it can just as easily stay the same or increase.
- Fast-DetectGPT (FastGPT) (Bao et al., 2024). To overcome the major computational expense of DetectGPT, this approach replaces DetectGPT's resource-intensive perturbation step with a more efficient sampling process. It identifies differences in token selection between humans and LLMs using a conditional probability curvature metric.
- **DNA-GPT** (Yang et al., 2024). This method involves a two-step process. First, it cuts a text in half and uses the first part to prompt an LLM to generate a new continuation. Next, it examines the differences between the newly created segment and the original one. This comparison, done via N-gram analysis for black-box models or probability divergence for white-box models, reveals a clear distinction between how humans and machines generate text.

D.3 EXPERIMENTAL SCENARIO

To extensively evaluate the effectiveness of the proposed enhancement model, we conduct experiments in the following real-world scenarios:

• Cross-LLM. To assess how well the proposed model works across different LLMs, we trained detectors on a single LLM's text and then tested it on a variety of LLMS' texts. The main body of the paper presents the results from training detectors on the GPT4All texts (Essay dataset) and Llama-2-70b texts (DetectRL), and then testing them on various LLMs, as shown in Table 2. Complete results for every training and testing combination can be found in Appendix E.3 (Tables 3-15).

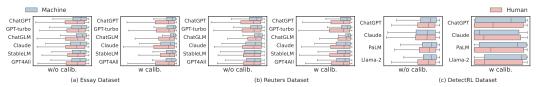


Figure 9: Distribution of token scores obtained by the Log-Likelihood method without and with enhancement. It can be observed that the proposed method enhances the discriminative nature of the token scores.

- Cross-Domain. The DetectRL dataset includes texts from four distinct domains: arXiv academic abstracts, XSum news articles, Writing Prompts stories, and Yelp Reviews. We utilized this dataset to evaluate the model's performance across various domains. To achieve this, we trained the detector on one domain and then tested it on the other domains. For this evaluation, all machine-generated texts were created using the default PaLM model. The heatmaps in Figs. 4, 19, and 20 illustrate the performance improvement achieved by our enhancement model compared to various baseline detection models.
- Paraphrasing Attack. Studies have shown that MGT detection is vulnerable to paraphrase attacks. Therefore, this scenario is used to evaluate the robustness of the MGT detection method. Using the DetectRL dataset, which includes data from the Polish and Dipper paraphraser, we trained our detector on clean, original texts and then evaluated its robustness on these paraphrased texts. Specifically, we trained the detector using clean texts from Llama-2-70b and then tested it on paraphrased texts from several different LLMs. The results can be found in Fig. 6 and 22.
- Mixed Text. Because a blend of human and machine-generated text is common in the real world, we use the mixed texts provided by DetectRL for evaluation. It involved randomly swapping out 25% of the sentences in an LLM-generated text with human-written ones. We conducted two separate experiments on this dataset: (1) The detector was trained on pure, non-mixed text and then tested for its ability to detect the mixed texts. (2) The detector was both trained and tested on the mixed texts themselves. The performance of the detector in these mixed settings is shown in Figs. 5 and 21. In each sub-figure, the detectors trained on original text are shown on the left, and those trained on mixed text are shown on the right.

D.4 PARAMETER SETTINGS

We conducted five independent experiments to ensure the consistency of our results, using five fixed random seeds (1-5). For all datasets, we used 10% of the data for training, while the remaining 90% was split evenly between validation and testing. To ensure a fair comparison, the enhanced models shared the same hyperparameters as their base models. In our enhancement model, there are two hyperparameters: the transition center t_0 in the positional weighting function $\beta(t)$ and the number of iterations T in the MRF layer. By default, we set $t_0=30$ and T=10 for the enhanced versions of all detectors across the three datasets, highlighting the flexibility of our approach. For training the MRF layer, we use a learning rate of 0.05 and train for 10 epochs. Hyperparameter sensitivity analyses are provided in Appendix E.5.

E MORE EXPERIMENTAL RESULS

E.1 More Results of Token Score Distribution before and after Enhancement

In addition to the partial results on DetectGPT presented in the main text, we also present the complete results about token-level detection score distributions for Log-Likelihood, Log-Rank, Entropy, and DetectGPT in Figs. 9, 10, 11, and 12. The results are similar to those in the main text: the original detector's scores show substantial overlap between human- and machine-generated text. However, after calibration using our proposed augmentation strategy, the scores achieve significantly improved discriminability.

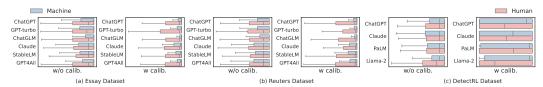


Figure 10: Distribution of token scores obtained by the Log-Rank method without and with enhancement. It can be observed that the proposed method enhances the discriminative nature of the token scores.

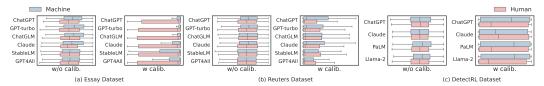


Figure 11: Distribution of token scores obtained by the Entropy method without and with enhancement. It can be observed that the proposed method enhances the discriminative nature of the token scores.

E.2 More Results of Context Token Relationships

In the main text, we demonstrated the existence of neighbor similarity (Fig. 2) and initial instability (Fig. 2) in token-level detection scores through experiments on the Essay dataset. Here, we provide additional supplementary results to further validate these relationships.

To verify neighbor similarity, we provide additional results using Entropy and DetectGPT detection scores on the Essay dataset (Fig. 13), as well as results on the Reuters and DetectRL datasets (Figs. 14 and 15). These supplementary experiments consistently show that the closer the tokens are, the more similar their detection scores are.

Similarly, to verify initial instability, we provide additional results on the Essay dataset (Fig. 16), as well as results on the Reuters and DetectRL datasets (Figs. 17 and 18). These results again demonstrate that token scores at the beginning of a text fluctuate significantly before gradually stabilizing.

E.3 More Performance Comparison

Performance across Different LLMs. In the main text, we evaluated the cross-LLM performance of detectors trained on GPT4All (Essay) and Llama-2-70b (DetectRL) on various LLM texts in terms of AUROC, as shown in Table 2. This section aims to provide more comprehensive supplementary experimental results, including: (1) cross-LLM performance performance under the same experimental settings in terms of TPR@FPR-1% (Table 3), and (2) cross-LLM performance comparisons of detectors trained on other LLMs on the Essay, DetectRL, and Reuters datasets (Tables 4 to 15). These extensive experimental results are consistent with the conclusions of the main paper. Among the 888 cross-LLM evaluation settings, our proposed enhanced model achieved better performance than the original detector in 91.4% of the cases, highlighting the generalization ability and application value of this method on different models and datasets.

Performance across Different Domains. In addition to the cross-domain performance improvements for Log-Likelihood and Log-Rank demonstrated in the main text, this section provides additional results for other detectors. Specifically, it includes improvements to the cross-domain performance of Entropy and DetectGPT (Fig. 19), as well as improvements to FastGPT and DNA-GPT (Fig. 20). Combining all experimental results, we reach the same conclusion as in the main text: in most experimental settings, detectors applied with our strategy significantly improve their cross-domain generalization capabilities.

Performance against Mixed Texts. In addition to the AUROC performance comparison for mixed texts presented in the main text, this section provides additional performance comparisons at TPR@FPR-1%, as shown in Fig. 21. We reach consistent conclusions with those in the main

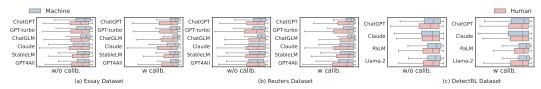


Figure 12: Distribution of token scores obtained by the DetectGPT method without and with enhancement. It can be observed that the proposed method enhances the discriminative nature of the token scores.

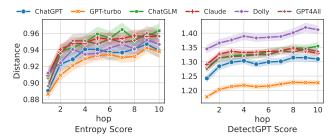


Figure 13: The detection score distances of neighbors at different hops in the Essay dataset. Entropy and DetectGPT score are used here.

text: in most experimental settings, the detector equipped with the proposed strategy significantly improves its ability to detect mixed text.

Performance against Paraphrasing Attacks. In addition to the AUROC performance comparison for paraphrasing texts presented in the main text, this section also presents a performance comparison under the TPR@FPR-1% metric, as shown in Fig. 22. We find that in most experimental settings, the detector applying our proposed strategy significantly improves its robustness against paraphrasing attacks, which is consistent with the conclusions drawn in the main text.

E.4 More Results of Ablation Study

In addition to the ablation experiments on the position weight function and MRF layer presented in the main text for the Essay dataset, we also provide ablation results for the Reuters and DetectRL datasets, as shown in Figs. 23 and 24. These experimental results are consistent with the conclusions of the main text: removing either component leads to a significant performance drop; even retaining only one of them outperforms the baseline detector, strongly demonstrating the effectiveness of both components.

E.5 SENSITIVITY ANALYSIS

Sensitivity w.r.t. transition center t_0 . In our experiments, we set the default transition center t_0 of the position weighting function to 30. This section examines the effect of varying t_0 values on detection performance. The AUROC and TPR@FPR-1% results are shown in Figs. 25 and 26, respectively. The experimental results show that detection performance steadily improves with increasing t_0 values, which is consistent with the conclusions of the ablation experiments and demonstrates the effectiveness of the position weighting function. However, performance improvement is not infinite. When t_0 values are too large, performance gradually saturates or even declines. This is likely because excessively large t_0 filters out useful token scores. Therefore, a trade-off is necessary. Through sensitivity analysis, we recommend setting t_0 values between 20 and 30.

Sensitivity w.r.t. iterations T in MRF layer. We compute the posterior probability of the Markov random field using a multi-step iterative approach. To this end, we evaluate the impact of varying the number of iterations on detection performance, as shown in Figs. 27 and 28. The experimental results demonstrate that multi-step iterative computation significantly enhances detection performance compared to single-step computation, underscoring its importance in score calibration. However, performance may degrade with increasing the number of iterations, possibly due to oversmoothing

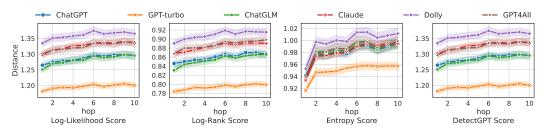


Figure 14: The detection score distances of neighbors at different hops in the Reuters dataset. Log-Likelihood, Log-Rank, Entropy, and DetectGPT score are used here.

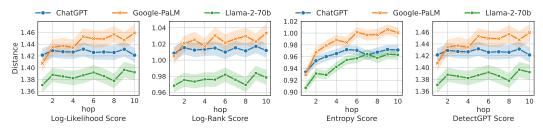


Figure 15: The detection score distances of neighbors at different hops in the DetectRL dataset. Log-Likelihood, Log-Rank, Entropy, and DetectGPT score are used here.

of the detection scores. Based on our sensitivity analysis, we recommend setting the number of iterations to 10.

E.6 More Results Comparing with Neural Network Calibration

In the main text, we demonstrated, through experimental results using the Log-Likelihood score, that while neural network-based detection score correction performs well on the intra-LLM, its performance drops sharply on the cross-LLM. This section further presents comparative results using other detection scores, as shown in Figs. 29 and 30. Similar experimental findings further demonstrate that this method does not truly learn universal score correction capabilities, but rather overfits the training data. This strongly emphasizes the superiority and rationality of our proposed Markovinformed score calibration method.

E.7 COMPARISON WITH MODEL-BASED DETECTORS

As shown in Fig. 31, we compare the enhanced versions of metric-based methods with model-based detection methods, including ChatGPT-D and MPU. Experimental results show that while model-based methods demonstrate superior performance on the DetectRL dataset, they underperform state-of-the-art metric-based methods, such as DNA-GPT, on the Essay and Reuters datasets. Notably, the significant performance gap between model-based methods in intra-LLM and cross-LLM further confirms their increased risk of overfitting to the training data. This observation is consistent with our intention of focusing on metric-based detection methods.

E.8 RUNNING TIME

Table 16 shows the training and inference runtimes on different datasets. As discussed in the main text, our proposed Markov-based score refinement module can be implemented in constant time via sparse-dense matrix multiplication. Therefore, the additional time overhead introduced by this module is negligible compared to the time-consuming score calculation, highlighting the flexibility and practicality of our approach in practical applications.

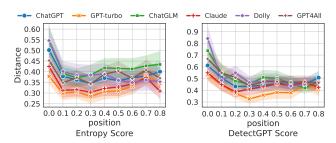


Figure 16: The score distances of 1-hop neighbors at different positions in the Essay dataset. Entropy and DetectGPT score are used here.

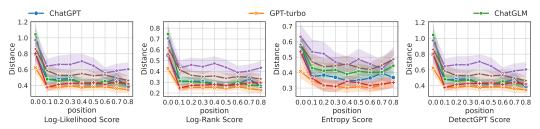


Figure 17: The score distances of 1-hop neighbors at different positions in the Reuters dataset. Log-Likelihood, Log-Rank, Entropy, and DetectGPT score are used here.

F THE USE OF LARGE LANGUAGE MODELS

In our paper, we used LLMs to polish the language and correct grammatical errors. LLMs were not used to generate novel research ideas, design experiments, analyze results, or write substantive technical content. We ensure that the use of large language models is responsible and adheres to academic and ethical standards.

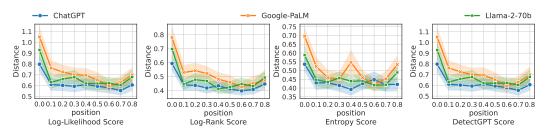


Figure 18: The score distances of 1-hop neighbors at different positions in the DetectRL dataset. Log-Likelihood, Log-Rank, Entropy, and DetectGPT score are used here.

Table 3: Performance concerning TPR@FPR-1% (%) on Essay (left) and DetectRL (right).

						` /		· /		· U /	
Method			Essay (Train	ing Text: GPT4	All)			DetectRL (Training Text: Llama-2-70b)			
	GPT4All	ChatGPT	Dolly	ChatGLM	Claude	ChatGPT-turbo	Avg.	Llama-2-70b	ChatGPT	Google-PaLM	Avg.
Likelihood	46.33±16.49	$68.62_{\pm 13.32}$	$20.67_{\pm 10.79}$	$92.86_{\pm 4.84}$	$12.31_{\pm 6.29}$	$73.60_{\pm 14.63}$	52.4	38.37+0.92	$10.21_{\pm 1.40}$	$25.66_{\pm 2.41}$	24.75
Likelihood-M	87.47 _{±3.42}	$90.36_{\pm 2.93}$	55.47 _{±5.03}	97.19 $_{\pm 1.11}$	$40.40_{\pm 7.87}$	$96.27_{\pm 1.27}$	77.86	46.03±1.70	$10.51_{\pm 3.92}$	$36.37_{\pm 2.31}$	30.97
Log-Rank	$62.69_{\pm 13.36}$	$79.07_{\pm 7.89}$	$25.11_{\pm 8.35}$	$95.71_{\pm 2.05}$	$19.20_{\pm 8.26}$	$80.89_{\pm 10.73}$	60.44	42.05 ± 0.84	$12.44_{\pm 1.71}$	$22.84_{\pm 1.28}$	25.78
Log-Rank-M	87.38 _{±2.50}	$88.89_{\pm 2.84}$	$52.12_{\pm 4.18}$	$98.04_{\pm0.43}$	$31.69_{\pm 4.15}$	94.22 _{±2.46}	75.39	50.56±1.44	$13.94_{\pm 2.36}$	$41.51_{\pm 2.69}$	35.34
Entropy	$2.73_{\pm 0.69}$	$7.16_{\pm 3.17}$	2.29 ± 1.36	$6.88_{\pm 3.05}$	$3.91_{\pm 2.24}$	$13.24_{\pm 7.60}$	6.04	$2.03_{\pm 0.73}$	$0.25_{\pm 0.16}$	$6.95_{\pm 0.78}$	3.07
Entropy-M	15.63 _{±0.87}	$41.24_{\pm 2.25}$	$16.95_{\pm 3.12}$	$40.85_{\pm 4.37}$	$23.20_{\pm 1.70}$	49.91 _{±5.02}	31.30	2.69 _{±0.39}	$0.67_{\pm 0.10}$	$9.17_{\pm 1.10}$	4.18
DetectGPT	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.89_{\pm0.24}$	$0.00_{\pm 0.00}$	0.15	$4.15_{\pm0.40}$	$2.94_{\pm 0.62}$	$1.71_{\pm 0.47}$	2.93
DetectGPT-M	41.23 _{±29.53}	46.98 ± 25.36	$18.38_{\pm 20.21}$	$41.83_{\pm 44.80}$	$8.80_{\pm 9.64}$	$65.87_{\pm 21.66}$	37.18	32.41 _{±9.58}	$7.59_{\pm 4.05}$	$25.14_{\pm 4.53}$	21.71
FastGPT	1.59 ± 0.32	$1.64_{\pm 0.67}$	$0.29_{\pm 0.10}$	2.72 ± 0.68	$3.47_{\pm 1.83}$	$16.89_{\pm 6.66}$	4.43	$11.35_{\pm 2.01}$	$3.78_{\pm 1.20}$	$5.02_{\pm 0.85}$	6.72
FastGPT-M	23.92 _{±10.29}	$30.00_{\pm 9.20}$	$32.89_{\pm 16.52}$	$55.54_{\pm 7.80}$	$0.67_{\pm 0.56}$	$16.44_{\pm 31.78}$	26.58	15.23±3.23	$6.28_{\pm 1.72}$	$13.52_{\pm 5.63}$	11.68
DNA-GPT	$57.68_{\pm 3.34}$	$63.24_{\pm 6.54}$	$22.20_{\pm 3.85}$	$84.82_{\pm 4.14}$	$16.67_{\pm 3.47}$	$56.27_{\pm 6.45}$	50.15	36.56 ± 1.56	19.70 $_{\pm 0.52}$	$30.38_{\pm 2.75}$	28.88
DNA-GPT-M	93.85 _{±0.97}	$88.44_{\pm 4.22}$	$56.42_{\pm 10.82}$	$96.83_{\pm0.41}$	$33.33_{\pm 7.90}$	$86.22_{\pm 6.52}$	75.85	42.52 _{±1.26}	$19.11_{\pm 2.04}$	$34.81_{\pm 1.54}$	32.15

Table 4: Performance on Essay dataset. The detection models are trained on text generated by ChatGPT.

Metric	Method	GPT4All	ChatGPT	Dolly	ChatGLM	Claude	ChatGPT-turbo	avg
	Likelihood	$46.24_{\pm 16.60}$	$68.62_{\pm 13.32}$	$20.67_{\pm 10.79}$	92.81 _{±4.82}	12.31 _{±6.29}	$73.60_{\pm 14.63}$	52.38
	Likelihood-M	89.38 +4.05	$92.00_{\pm 2.84}$	$59.19_{\pm 5.98}$	$97.10_{\pm 0.69}$	$42.22_{\pm 6.32}$	$96.80_{\pm 1.23}$	79.45
	Log-Rank	$62.69_{\pm 13.36}$	$79.07_{\pm 7.89}$	$25.11_{\pm 8.35}$	$95.71_{\pm 2.05}$	19.16 ± 8.22	$80.89_{\pm 10.73}$	60.44
	Log-Rank-M	87.15+2.62	$88.80_{\pm 2.81}$	$52.12_{\pm 4.32}$	$98.04_{\pm0.43}$	$31.33_{\pm 4.10}$	$94.22_{+2.46}$	75.28
	Entropy	$2.73_{\pm 0.69}^{\pm 2.02}$	$7.16_{\pm 3.17}$	$2.29_{\pm 1.36}$	$6.88_{\pm 3.05}$	$3.91_{\pm 2.24}$	$13.24_{\pm 7.60}$	6.04
ΓPR@FPR-1%	Entropy-M	$12.76_{\pm 2.18}$	$34.84_{\pm 5.15}$	$14.27_{\pm 2.63}$	$36.96_{\pm 2.81}$	19.78 _{+3.76}	$43.20_{\pm 6.33}$	26.97
irkerrk-r/	DetectGPT	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.84_{\pm 0.51}$	2.36 ± 2.52	0.53
	DetectGPT-M	$41.09_{\pm 30.67}$	$42.04_{\pm 25.95}$	$9.55_{\pm 9.01}$	$39.02_{\pm 38.26}$	$10.13_{\pm 10.28}$	$68.49_{+23.15}$	35.05
	FastGPT	$1.59_{\pm 0.32}$	$1.64_{\pm 0.67}$	$0.29_{\pm 0.10}$	$2.72_{\pm 0.68}$	$3.47_{\pm 1.83}$	16.89 ± 6.66	4.43
	FastGPT-M	$19.73_{+11.16}$	$26.18_{\pm 12.42}$	$24.82_{\pm 20.23}$	$45.76_{+21.65}$	$3.38_{\pm 5.34}$	$26.04_{\pm 32.58}$	24.32
	DNA-GPT	$40.77_{\pm 8.33}$	$73.69_{\pm 3.78}$	$19.95_{\pm 7.42}$	$82.59_{\pm 11.71}$	$18.04_{\pm 3.53}$	$63.82_{\pm 6.62}$	49.81
	DNA-GPT-M	$89.02_{\pm 4.14}$	95.07 \pm 1.33	$65.97_{\pm 4.57}$	99.24 ±0.36	$39.56_{\pm 4.01}$	$94.62_{\pm 1.43}$	80.58
	Likelihood	$96.16_{\pm0.30}$	$98.79_{\pm0.19}$	$90.90_{\pm 1.33}$	$99.29_{\pm 0.25}$	$92.76_{\pm0.23}$	99.13 _{±0.19}	96.17
	Likelihood-M	$98.77_{\pm 0.10}$	99.58 \pm 0.14	94.80 ± 0.87	99.56+0.19	$94.69_{\pm 0.42}$	99.77 $_{\pm 0.14}$	97.86
	Log-Rank	$96.55_{\pm0.31}$	98.95 ± 0.13	90.08 ± 1.28	$99.36_{\pm0.13}$	$92.01_{\pm 0.20}$	$99.24_{\pm 0.15}$	96.03
	Log-Rank-M	98.56 _{+0.06}	$99.40_{\pm 0.09}$	$93.81_{\pm 0.92}$	99.56 _{+0.08}	$92.88_{\pm 0.29}$	99.64 _{+0.10}	97.31
	Entropy	$74.19_{\pm 1.62}$	$89.49_{\pm 0.33}$	$73.26_{\pm 1.48}$	$84.11_{\pm 0.77}$	$86.58_{\pm 0.66}$	$95.94_{\pm 0.35}$	83.93
AUROC	Entropy-M	$83.33_{\pm 0.66}$	$93.24_{\pm 0.15}$	$81.17_{\pm 1.04}$	$91.38_{\pm0.37}$	$88.28_{\pm0.49}$	$96.86_{\pm0.21}$	89.04
AUROC	DetectGPT	50.15 ± 0.99	50.53 ± 3.65	48.04 ± 7.27	$49.25_{\pm 1.58}$	$51.21_{\pm 8.39}$	55.34 ± 31.67	50.75
	DetectGPT-M	94.98 _{+3.26}	96.35 ± 2.60	$85.82_{\pm 6.13}$	$94.20_{\pm 4.48}$	84.06 ± 6.47	$98.58_{\pm 1.24}$	92.33
	FastGPT	$64.63_{\pm 1.53}$	$67.68_{\pm 1.70}$	$47.17_{\pm 1.53}$	$71.08_{\pm 1.51}$	$75.31_{\pm 0.90}$	$88.62_{\pm 0.67}$	69.08
	FastGPT-M	86.68+3.28	$90.82_{\pm 3.86}$	$72.23_{+23.91}$	$93.53_{\pm 3.86}$	$65.24_{\pm 12.36}$	$71.87_{\pm 21.80}$	80.06
	DNA-GPT	$96.28_{\pm0.32}$	$98.87_{\pm0.21}$	$92.85_{\pm0.83}$	$99.32_{\pm 0.27}$	$91.65_{\pm 0.72}$	$98.51_{\pm 0.31}$	96.25
	DNA-GPT-M	99.32 \pm 0.13	99.79 \pm 0.07	97.42 $_{\pm 0.56}$	99.90±0.06	95.65 ± 0.51	99.76 ±0.07	98.64

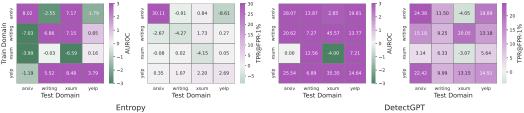


Figure 19: The performance improvement of the proposed method on Entropy and DetectGPT. Values greater than 0 indicate an enhanced effect.

Table 5: Performance on Essay dataset. The detection models are trained on text generated by ChatGPT-turbo.

Metric	Method	GPT4All	ChatGPT	Dolly	ChatGLM	Claude	ChatGPT-turbo	avg
	Likelihood	46.29 _{±16.55}	$68.62_{\pm 13.32}$	$20.67_{\pm 10.79}$	$92.86_{\pm 4.84}$	$12.31_{\pm 6.29}$	$73.60_{\pm 14.63}$	52.39
	Likelihood-M	87.61 ± 3.74	88.31 ± 3.64	61.29 \pm 7.12	$94.20_{\pm 2.40}$	$41.11_{\pm 7.93}$	$95.38_{\pm 2.27}$	77.98
	Log-Rank	$62.69_{\pm 13.36}$	$79.07_{\pm 7.89}$	$25.11_{\pm 8.35}$	$95.71_{\pm 2.05}$	19.16 ± 8.22	80.89 ± 10.73	60.44
	Log-Rank-M	86.70 ±0.59	$87.47_{\pm 3.17}$	$54.61_{\pm 4.05}$	93.75 ± 3.55	$35.91_{\pm 2.65}$	$94.13_{\pm 1.68}$	75.43
	Entropy	$2.73_{\pm 0.69}$	$7.16_{\pm 3.17}$	$2.29_{\pm 1.36}$	$6.88_{\pm 3.05}$	$3.91_{\pm 2.24}$	$13.24_{\pm 7.60}$	6.04
TPR@FPR-1%	Entropy-M	11.62 ± 0.59	30.67 ± 6.02	13.56 ± 1.59	$32.54_{\pm 2.63}$	$17.33_{\pm 5.32}$	$39.91_{\pm 8.11}$	24.27
11 K@11 K-1 //	DetectGPT	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	0.53 ± 0.52	$3.07_{\pm 1.96}$	0.60
	DetectGPT-M	$14.03_{\pm 13.72}$	$17.07_{\pm 8.79}$	$3.05_{\pm 2.54}$	$6.43_{\pm 3.80}$	$5.11_{\pm 3.98}$	$44.71_{\pm 11.72}$	15.07
	FastGPT	$1.59_{\pm 0.32}$	$1.64_{\pm 0.67}$	$0.29_{\pm 0.10}$	$2.72_{\pm 0.68}$	$3.47_{\pm 1.83}$	$16.89_{\pm 6.66}$	4.43
	FastGPT-M	$8.97_{\pm 4.52}$	$15.33_{\pm 15.63}$	$0.72_{\pm 0.60}$	$20.09_{\pm 18.21}$	$7.51_{\pm 6.83}$	$49.82_{\pm 22.08}$	17.07
	DNA-GPT	$0.73_{\pm 0.33}$	$0.76_{\pm 0.30}$	$2.63_{\pm 0.81}$	$0.13_{\pm 0.11}$	$21.64_{\pm 3.45}$	$0.00_{\pm 0.00}$	4.31
	DNA-GPT-M	$39.41_{\pm 2.68}$	$38.62_{\pm 2.34}$	$39.47_{\pm 2.11}$	$38.04_{\pm 2.49}$	$39.38_{\pm 1.86}$	$38.67_{\pm 2.19}$	38.93
	Likelihood	96.16 _{±0.30}	$98.79_{\pm0.19}$	$90.90_{\pm 1.33}$	99.29 _{±0.25}	$92.76_{\pm0.23}$	$99.13_{\pm0.19}$	96.17
	Likelihood-M	$98.65_{\pm 0.21}$	$99.39_{\pm 0.20}$	$93.95_{\pm 1.15}$	$99.13_{\pm 0.43}$	94.66 \pm 0.39	99.77 $_{\pm 0.12}$	97.59
	Log-Rank	$96.55_{\pm0.31}$	$98.95_{\pm0.13}$	$90.08_{\pm 1.28}$	99.36 $_{\pm 0.13}$	$92.01_{\pm 0.20}$	$99.24_{\pm 0.15}$	96.03
	Log-Rank-M	$98.37_{\pm 0.21}$	99.17 _{+0.29}	$92.96_{\pm 1.51}$	$99.17_{\pm 0.37}$	$92.96_{\pm 0.43}$	99.62 $_{\pm 0.11}$	97.04
	Entropy	$74.19_{\pm 1.62}$	$89.49_{\pm 0.34}$	$73.26_{\pm 1.48}$	$84.11_{\pm 0.77}$	$86.58_{\pm 0.66}$	$95.94_{\pm 0.35}$	83.93
AUROC	Entropy-M	82.88 _{+0.85}	$93.08_{\pm 0.29}$	$80.90_{\pm 0.89}$	$91.18_{\pm 0.37}$	$88.57_{\pm 0.54}$	$96.90_{\pm0.24}$	88.92
AUROC	DetectGPT	$ 49.19_{\pm 0.58} $	$53.60_{\pm 0.77}$	$42.52_{\pm 0.84}$	$49.59_{\pm 1.70}$	$58.46_{\pm 0.60}$	$82.10_{\pm 1.25}$	55.91
	DetectGPT-M	91.65 _{+2.28}	$93.38_{+2.31}$	77.58 ± 5.63	$90.54_{\pm 2.74}$	79.82 $_{\pm 9.52}$	$97.61_{\pm 0.92}$	88.43
	FastGPT	$64.63_{\pm 1.53}$	$67.68_{\pm 1.70}$	$47.17_{\pm 1.53}$	$71.08_{\pm 1.51}$	$75.31_{\pm 0.90}$	$88.62_{\pm 0.67}$	69.08
	FastGPT-M	85.05 _{+3.63}	$86.41_{\pm 10.08}$	$49.85_{\pm 11.80}$	$89.08_{\pm 6.67}$	$75.84_{\pm 17.10}$	$96.80_{\pm 2.60}$	80.51
	DNA-GPT	$53.28_{\pm 1.23}$	$43.65_{\pm 0.76}$	$61.90_{\pm 0.39}$	$40.46_{\pm 1.26}$	$67.53_{\pm 0.55}$	$43.31_{\pm 0.94}$	51.69
	DNA-GPT-M	$83.94_{\pm 1.09}$	$84.02_{\pm 1.06}$	$83.93_{\pm 1.20}$	$83.89_{\pm 1.10}^{-}$	$84.23_{\pm 1.07}$	84.07 $_{\pm 1.06}$	84.01

Table 6: Performance on Essay dataset. The detection models are trained on text generated by ChatGLM.

Metric	Method	GPT4All	ChatGPT	Dolly	ChatGLM	Claude	ChatGPT-turbo	avg
TPR@FPR-1%	Likelihood Likelihood-M Log-Rank Log-Rank-M Entropy-M DetectGPT-M FastGPT-M FastGPT-M DNA-GPT-M	$\begin{array}{c} \textbf{46}.29 \pm 16.54 \\ \textbf{78}.82 \pm 5.30 \\ \textbf{62}.69 \pm 13.36 \\ \textbf{81}.32 \pm 7.05 \\ \textbf{2}.73 \pm 0.69 \\ \textbf{12}.76 \pm 1.21 \\ \textbf{0}.00 \pm 0.00 \\ \textbf{1}.59 \pm 0.32 \\ \textbf{23}.87 \pm 10.38 \\ \textbf{57}.54 \pm 4.18 \\ \textbf{93}.58 \pm 1.29 \end{array}$	88.62±13.32 87.87±3.22 79.07±7.89 86.76±3.68 7.16±3.17 32.09±4.82 0.00±0.00 17.87±5.87 1.64±0.67 29.87±9.03 69.91±9.92 92.53±3.44	$\begin{array}{c} 20.67 \pm 10.79 \\ \textbf{46.73} \pm 10.16 \\ 25.11 \pm 8.35 \\ \textbf{47.35} \pm 8.02 \\ 2.29 \pm 1.36 \\ \textbf{13.99} \pm 3.01 \\ 0.00 \pm 0.00 \\ \textbf{20.67} \pm 3.81 \\ 0.29 \pm 0.10 \\ \textbf{32.89} \pm 16.52 \\ 39.19 \pm 4.23 \\ \textbf{66.54} \pm 9.55 \end{array}$	$\begin{array}{c} 92.81 \pm 4.82 \\ \textbf{98.71} \pm 1.11 \\ 95.71 \pm 2.05 \\ \textbf{98.62} \pm 0.26 \\ 6.88 \pm 3.05 \\ \textbf{35.62} \pm 4.65 \\ 0.00 \pm 0.00 \\ \textbf{60.98} \pm 6.42 \\ 2.72 \pm 0.68 \\ \textbf{55.49} \pm 7.87 \\ \textbf{97.32} \pm 0.76 \\ \textbf{99.55} \pm 0.20 \end{array}$	$\begin{array}{c} 12.31 \pm 6.29 \\ \textbf{31.64} \pm 10.61 \\ 19.20 \pm 8.26 \\ \textbf{25.24} \pm 7.46 \\ 3.91 \pm 2.24 \\ \textbf{18.40} \pm 4.10 \\ 0.89 \pm 0.24 \\ \textbf{4.89} \pm 2.75 \\ \textbf{3.47} \pm 1.83 \\ 0.67 \pm 0.56 \\ \textbf{31.24} \pm 5.75 \\ \textbf{44.36} \pm 9.19 \end{array}$	$\begin{array}{c} \textbf{73.60} \pm 14.63 \\ \textbf{93.78} \pm 3.41 \\ \textbf{80.89} \pm 10.73 \\ \textbf{90.62} \pm 5.22 \\ \textbf{13.24} \pm 7.60 \\ \textbf{41.42} \pm 6.98 \\ \textbf{0.00} \pm 0.00 \\ \textbf{5.06} \pm 2.01 \\ \textbf{16.89} \pm 6.66 \\ \textbf{16.22} \pm 31.34 \\ \textbf{65.16} \pm 10.11 \\ \textbf{91.47} \pm 4.41 \end{array}$	52.38 72.92 60.44 71.65 6.04 25.71 0.15 21.52 4.43 26.50 60.06 81.34
AUROC	Likelihood Likelihood-M Log-Rank Log-Rank-M Entropy-M DetectGPT DetectGPT-M FastGPT-M DNA-GPT DNA-GPT-M	$\begin{array}{c} 96.16 \pm 0.30 \\ 98.03 \pm 0.26 \\ 96.55 \pm 0.31 \\ 98.14 \pm 0.30 \\ 74.19 \pm 1.62 \\ 82.99 \pm 1.01 \\ 50.81 \pm 0.58 \\ 80.39 \pm 1.33 \\ 64.63 \pm 1.53 \\ 87.19 \pm 3.46 \\ 96.97 \pm 0.20 \\ 99.37 \pm 0.05 \\ \end{array}$	$\begin{array}{c} 98.79 \pm 0.19 \\ 99.22 \pm 0.20 \\ 98.95 \pm 0.13 \\ 99.20 \pm 0.22 \\ 89.49 \pm 0.33 \\ 93.13 \pm 0.23 \\ 46.40 \pm 0.77 \\ 80.53 \pm 0.58 \\ 67.68 \pm 1.70 \\ 91.55 \pm 3.32 \\ 98.09 \pm 0.47 \\ 99.58 \pm 0.10 \\ \end{array}$	$\begin{array}{c} 90.90\pm1.33\\ \textbf{93.10}\pm1.43\\ 90.08\pm1.28\\ \textbf{92.86}\pm1.66\\ 73.26\pm1.48\\ \textbf{81.03}\pm1.11\\ 57.48\pm0.84\\ \textbf{80.67}\pm1.98\\ 47.17\pm1.53\\ \textbf{82.58}\pm18.04\\ 94.11\pm0.47\\ \textbf{97.68}\pm0.29\\ \end{array}$	$\begin{array}{c} 99.29 \pm 0.25 \\ \textbf{99.65} \pm 0.07 \\ 99.36 \pm 0.13 \\ \textbf{99.54} \pm 0.10 \\ 84.11 \pm 0.77 \\ \textbf{91.15} \pm 0.28 \\ 50.41 \pm 1.70 \\ \textbf{94.02} \pm 0.56 \\ 71.08 \pm 1.51 \\ \textbf{95.35} \pm 0.49 \\ \textbf{99.87} \pm 0.04 \\ \textbf{99.92} \pm 0.02 \\ \end{array}$	$\begin{array}{c} 92.76 \pm 0.23 \\ 93.47 \pm 1.14 \\ 92.01 \pm 0.20 \\ 91.39 \pm 1.08 \\ 86.58 \pm 0.66 \\ 88.45 \pm 0.44 \\ 41.54 \pm 0.60 \\ \textbf{62.63} \pm 1.24 \\ \textbf{75.31} \pm 0.90 \\ 59.28 \pm 1.90 \\ 92.84 \pm 0.71 \\ \textbf{96.24} \pm 0.47 \end{array}$	$\begin{array}{c} 99.13 \pm 0.19 \\ 99.54 \pm 0.16 \\ 99.24 \pm 0.15 \\ 99.53 \pm 0.16 \\ 95.94 \pm 0.35 \\ 96.88 \pm 0.18 \\ 17.90 \pm 1.25 \\ 67.25 \pm 1.24 \\ 88.62 \pm 0.67 \\ 63.48 \pm 18.12 \\ 97.85 \pm 0.37 \\ 99.54 \pm 0.09 \\ \end{array}$	96.17 97.17 96.03 96.78 83.93 88.94 44.09 77.58 69.08 79.91 96.62 98.72

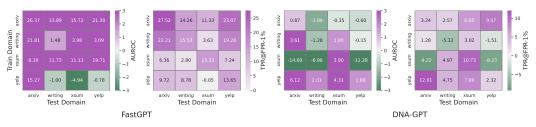


Figure 20: The performance improvement of the proposed method on FastGPT and DNA-GPT. Values greater than 0 indicate an enhanced effect.

Table 7: Performance on Essay dataset. The detection models are trained on text generated by Dolly.

Metric	Method	GPT4All	ChatGPT	Dolly	ChatGLM	Claude	ChatGPT-turbo	
	Likelihood Likelihood-M	$ 46.29_{\pm 16.54} $ 89.02 $_{\pm 1.68}$	$68.62_{\pm 13.32}$ $88.09_{\pm 4.50}$	$20.67_{\pm 10.79}$ 57.14 $_{\pm 5.55}$	$92.86_{\pm 4.84}$ $97.59_{\pm 1.09}$	$12.31_{\pm 6.29}$ 34.36 $_{\pm 13.54}$	$73.60_{\pm 14.63}$ $94.36_{\pm 3.67}$	52.39 76.76
	Log-Rank	$62.69_{\pm 13.36}$	$79.07_{\pm 7.89}$	25.11 ± 8.35	$95.71_{\pm 2.05}$	19.16 ± 8.22	$80.89_{\pm 10.73}$	60.44
	Log-Rank-M	87.70 ± 2.60	$89.02_{\pm 2.13}$	$51.65_{\pm 3.99}$	98.12 \pm 0.54	$31.69_{\pm 3.32}$	$94.36_{\pm 1.93}$	75.42
	Entropy	2.73 ± 0.69	7.16 ± 3.17	$2.29_{\pm 1.36}$	6.88 ± 3.05	$3.91_{\pm 2.24}$	$13.24_{\pm 7.60}$	6.04
TPR@FPR-1%	Entropy-M	11.94 \pm 1.85	$32.09_{\pm 4.77}$	$13.22_{\pm 4.31}$	$34.78_{\pm 8.04}$	18.49 \pm 3.20	40.62 ± 6.34	25.19
IFK@IFK-1%	DetectGPT	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.89_{\pm0.24}$	$0.00_{\pm 0.00}$	0.15
	DetectGPT-M	$32.89_{\pm 22.87}$	$46.09_{\pm 19.75}$	$24.44_{\pm 9.98}$	$68.84_{\pm 21.48}$	$6.04_{\pm 4.69}$	$48.09_{\pm 25.20}$	37.73
	FastGPT	$0.82_{\pm 0.31}$	$0.00_{\pm 0.00}$	$0.38_{\pm0.42}$	$0.09_{\pm 0.11}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	0.22
	FastGPT-M	$27.24_{\pm 2.77}$	$28.04_{\pm 7.40}$	$39.43_{\pm 3.59}$	$55.89_{\pm 5.30}$	$0.98_{\pm 0.65}$	$0.84_{\pm0.71}$	25.41
j.	DNA-GPT	$58.95_{\pm 4.76}$	$64.00_{\pm 6.45}$	$48.64_{\pm 4.04}$	$84.29_{\pm 3.64}$	$29.24_{\pm 3.40}$	$59.02_{\pm 8.29}$	57.36
	DNA-GPT-M	$88.70_{\pm 1.78}$	$87.82_{\pm 1.27}$	$72.60_{\pm 4.17}$	$92.59_{\pm 0.38}$	$44.71_{\pm 4.67}$	$86.93_{\pm 2.13}$	78.89
	Likelihood	96.16 _{±0.30}	$98.79_{\pm0.19}$	$90.90_{\pm 1.33}$	$99.29_{\pm 0.25}$	$92.76_{\pm0.24}$	99.13 _{±0.19}	96.17
	Likelihood-M	98.63 _{±0.29}	99.48 $_{\pm 0.16}$	$94.44_{\pm 0.71}$	$99.62_{\pm 0.08}$	$94.49_{\pm 0.73}$	$99.70_{\pm 0.15}$	97.73
	Log-Rank	$96.55_{\pm 0.31}$	$98.95_{\pm0.13}$	$90.08_{\pm 1.28}$	$99.36_{\pm0.13}$	$92.01_{\pm 0.20}$	$99.24_{\pm 0.15}$	96.03
	Log-Rank-M	$98.55_{\pm 0.08}$	99.40 $_{\pm 0.09}$	$93.78_{\pm 0.93}$	99.56 \pm 0.09	$92.81_{\pm 0.38}$	99.65 $_{\pm 0.10}$	97.29
	Entropy	$ 74.19_{\pm 1.62} $	$89.49_{\pm0.33}$	$73.26_{\pm 1.48}$	$84.11_{\pm 0.77}$	$86.58_{\pm0.66}$	$95.94_{\pm0.35}$	83.93
AUROC	Entropy-M	$83.09_{\pm0.84}$	$93.15_{\pm 0.17}$	$81.05_{\pm 1.14}$	91.26 $_{\pm 0.24}$	$88.42_{\pm 0.52}$	$96.88_{\pm 0.20}$	88.97
AUROC	DetectGPT	$50.81_{\pm 0.58}$	$46.40_{\pm 0.77}$	$57.48_{\pm 0.84}$	$50.41_{\pm 1.70}$	$41.54_{\pm0.60}$	$17.90_{\pm 1.25}$	44.09
	DetectGPT-M	$93.62_{\pm 5.70}$	$93.95_{\pm 6.36}$	$90.01_{\pm 3.69}$	$97.82_{\pm 1.98}$	$81.42_{\pm 9.59}$	$92.20_{\pm 11.34}$	91.50
	FastGPT	$35.37_{\pm 1.53}$	$32.32_{\pm 1.70}$	$52.83_{\pm 1.53}$	$28.92_{\pm 1.51}$	$24.69_{\pm 0.90}$	$11.38_{\pm 0.67}$	30.92
	FastGPT-M	$88.35_{\pm0.89}$	89.67 $_{\pm 1.11}$	$91.55_{\pm 0.28}$	$95.27_{\pm 0.38}$	$59.65_{\pm 1.54}$	$54.61_{\pm 2.30}$	79.85
	DNA-GPT	$94.50_{\pm 0.07}$	$93.95_{\pm0.19}$	$95.91_{\pm 0.49}$	$95.04_{\pm0.15}$	$89.35_{\pm0.39}$	$93.81_{\pm 0.20}$	93.76
	DNA-GPT-M	$98.23_{\pm0.21}$	97.67 $_{\pm 0.16}$	$98.01_{\pm0.33}$	$98.00_{\pm0.14}$	$94.63_{\pm0.40}$	$97.64_{\pm0.17}$	97.37

Table 8: Performance on Essay dataset. The detection models are trained on text generated by Claude

Claude.								
Metric	Method	GPT4All	ChatGPT	Dolly	ChatGLM	Claude	ChatGPT-turbo	avg
	Likelihood Likelihood-M	$46.20_{\pm 16.47}$ 86.74 $_{\pm 3.37}$	$68.58_{\pm 13.28}$ $90.09_{\pm 2.80}$	$20.67_{\pm 10.79}$ 50.07 _{± 9.80}	$92.86_{\pm 4.84}$ $98.08_{\pm 0.27}$	$12.31_{\pm 6.29}$ $34.00_{\pm 6.66}$	$73.60_{\pm 14.63}$ 95.78 $_{\pm 1.40}$	52.37 75.79
	Log-Rank	$62.69_{\pm 13.36}$	$79.07_{\pm 7.89}$	$25.11_{\pm 8.35}$	$95.71_{\pm 2.05}$	$19.20_{\pm 8.26}$	$80.89_{\pm 10.73}$	60.44
	Log-Rank-M	$89.20_{\pm 2.33}$	$88.18_{\pm 2.58}$	$54.42_{\pm 3.96}$	$95.00_{\pm 3.56}$	$37.82_{\pm 2.28}$	$94.80_{\pm 1.37}$	76.57
	Entropy	$2.73_{\pm 0.69}$	$7.16_{\pm 3.17}$	$2.29_{\pm 1.36}$	$6.88_{\pm 3.05}$	$3.91_{\pm 2.24}$	$13.24_{\pm 7.60}$	6.04
TPR@FPR-1%	Entropy-M	$12.16_{\pm0.86}$	$29.07_{\pm 3.22}$	10.74 $_{\pm 4.12}$	$30.94_{\pm 5.62}$	$16.80_{\pm 2.73}$	$39.69_{\pm 4.99}$	23.23
11 K@11 K-1 //	DetectGPT	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}^{-}$	$0.00_{\pm 0.00}$	$0.84_{\pm 0.51}$	$2.36_{\pm 2.52}$	0.53
	DetectGPT-M	$34.62_{\pm 17.61}$	$38.71_{\pm 23.27}$	$24.73_{\pm 8.38}$	74.06 \pm 17.99	$3.91_{\pm 2.05}$	29.11 ± 24.44	34.19
	FastGPT	1.59 ± 0.32	1.64 ± 0.67	$0.29_{\pm 0.10}$	2.72 ± 0.68	$3.47_{\pm 1.83}$	$16.89_{\pm 6.66}$	4.43
	FastGPT-M	$27.47_{\pm 2.92}$	$28.04_{\pm 7.49}$	$39.57_{\pm 3.76}$	56.47 \pm 6.13	0.98 ± 0.65	$0.84_{\pm 0.71}$	25.56
	DNA-GPT	$0.50_{\pm 0.17}$	$0.31_{\pm 0.11}$	$0.67_{\pm 0.55}$	$0.00_{\pm 0.00}$	$0.13_{\pm 0.11}$	0.53 ± 0.36	0.36
	DNA-GPT-M	12.76 ± 2.00	$13.16_{\pm 2.36}$	$14.03_{\pm 3.12}$	12.14 \pm 1.73	$13.69_{\pm 2.59}$	$13.33_{\pm 2.38}$	13.19
-	Likelihood	96.16+0.30	$98.79_{\pm0.19}$	90.90 _{±1.33}	$99.29_{\pm 0.25}$	$92.76_{\pm0.23}$	99.13+0.19	96.17
	Likelihood-M	$98.52_{\pm 0.24}$	99.45 _{±0.09}	$94.32_{\pm 0.95}$	99.63 _{±0.09}	$94.35_{\pm 0.42}$	99.66 _{±0.16}	97.66
	Log-Rank	$96.55_{\pm 0.31}$	$98.95_{\pm 0.13}$	$90.08_{\pm 1.28}$	99.36 \pm 0.09	$92.01_{\pm 0.20}$	$99.24_{\pm 0.15}$	96.03
	Log-Rank-M	$98.47_{\pm 0.29}$	$99.28_{\pm 0.22}$	$93.15_{\pm 1.37}$	$99.25_{\pm 0.32}$	$92.93_{\pm 0.43}$	99.66 ± 0.10	97.13
	Entropy	$74.19_{\pm 1.62}$	$89.49_{\pm 0.33}$	$73.26_{\pm 1.48}$	$84.11_{\pm 0.77}$	$86.58_{\pm0.66}$	$95.94_{\pm0.35}$	83.93
ALIBOC	Entropy-M	$82.36_{\pm 1.15}$	$92.90_{\pm 0.26}$	$80.64_{\pm 1.12}$	$90.88_{\pm 0.52}$	$88.62_{\pm 0.48}$	$96.91_{\pm 0.23}$	88.72
AUROC	DetectGPT	$50.15_{\pm 0.99}$	$50.53_{\pm 3.65}$	$48.04_{\pm 7.27}$	$49.25_{\pm 1.58}$	$51.21_{\pm 8.39}$	$55.34_{\pm 31.67}$	50.75
	DetectGPT-M	86.96 _{±7.59}	$87.54_{\pm 7.67}$	$85.22_{\pm 5.50}$	$96.46_{\pm 2.24}$	$71.40_{+9.73}$	$80.33_{\pm 14.18}$	84.65
	FastGPT	$64.63_{\pm 1.53}$	$67.68_{\pm 1.70}$	$47.17_{\pm 1.53}$	$71.08_{\pm 1.51}$	$75.31_{\pm 0.90}$	$88.62_{\pm 0.67}$	69.08
	FastGPT-M	$88.45_{\pm 0.98}$	89.73 $_{\pm 1.25}$	$91.61_{\pm 0.36}$	$95.35_{\pm 0.47}$	$59.77_{\pm 1.62}$	$54.76_{\pm 2.39}$	79.95
	DNA-GPT	45.85 ± 2.15	$39.61_{\pm 1.52}$	$53.93_{\pm 1.47}$	$36.19_{\pm 1.75}$	$53.01_{\pm 1.67}$	40.65 ± 1.60	44.87
	DNA-GPT-M	76.26 ±1.36	76.25 \pm 1.29	76.38 \pm 1.33	75.98 \pm 1.39	76.74 ± 1.26	76.30 $_{\pm 1.29}$	76.32

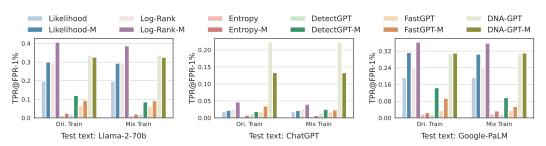


Figure 21: Detection performance concerning TPR@FPR-1% under different LLM mixed texts. All detectors are trained on Llama-2-70b texts.

Table 9: Performance on DetectRL. The detection models are trained on text generated by ChatGPT and Google-PaLM.

Metric	Method	Detec	tRL (Training	Text: ChatGPT)		DetectR	L (Training Tex	t: Google-PaLM	(I)
		Llama-2-70b	ChatGPT	Google-PaLM	Avg.	Llama-2-70b	ChatGPT	Google-PaLM	Avg.
	Likelihood	38.37±0.92	$10.21_{\pm 1.40}$	$25.66_{\pm 2.41}$	24.75	38.32±0.88	10.19 _{±1.38}	$25.66_{\pm 2.41}$	24.72
	Likelihood-M	$45.61_{\pm 3.57}$	$12.14_{\pm 2.66}$	$33.18_{\pm 2.46}$	30.31	$42.94_{\pm 4.14}$	$9.91_{\pm 3.87}$	$35.48_{\pm 2.70}$	29.44
	Log-Rank	42.05+0.84	$12.44_{\pm 1.71}$	$22.84_{\pm 1.28}$	25.78	42.05 ± 0.84	12.41 + 1.72	$22.84_{\pm 1.28}$	25.77
	Log-Rank-M	$50.66_{\pm 1.43}$	12.73 $_{\pm 1.65}$	$41.41_{\pm 2.56}$	34.93	50.58 _{+1.35}	$13.00_{\pm 2.43}$	$41.09_{\pm 4.07}$	34.89
	Entropy	2.03 ± 0.73	$0.25_{\pm 0.16}$	$6.95_{\pm 0.78}$	3.07	$2.03_{\pm 0.73}$	$0.25_{\pm 0.16}$	$6.95_{\pm 0.78}$	3.07
TPR@FPR-1%	Entropy-M	$2.13_{\pm 0.30}$	$0.27_{\pm 0.20}$	$9.27_{\pm 0.87}$	3.89	2.22 _{±0.37}	$0.27_{\pm 0.20}$	$9.37_{\pm 0.93}$	3.96
TPR@FPR-1%	DetectGPT	$3.29_{\pm 0.97}$	$5.36_{\pm 2.20}$	$2.79_{\pm 1.32}$	3.82	$2.87_{\pm 0.56}$	$6.77_{\pm 1.50}$	$3.86_{\pm 1.44}$	4.50
	DetectGPT-M	$25.29_{\pm 10.05}$	$7.66_{\pm 4.17}$	$14.17_{\pm 11.10}$	15.71	$22.00_{\pm 4.30}$	$4.57_{\pm 0.37}$	$24.94_{\pm 2.68}$	17.17
	FastGPT	$11.35_{\pm 2.01}$	$3.78_{\pm 1.20}$	$5.02_{\pm 0.85}$	6.72	$11.35_{\pm 2.01}$	$3.78_{\pm 1.20}$	$5.02_{\pm 0.85}$	6.72
	FastGPT-M	27.71 _{±15.63}	$13.42_{\pm 8.91}$	17.68 + 11.92	19.60	$12.41_{\pm 0.89}$	6.18 ± 0.83	$20.10_{\pm 1.78}$	12.90
	DNA-GPT	$31.03_{\pm 1.59}$	$11.05_{\pm 1.69}$	$26.40_{\pm 2.08}$	22.83	$41.71_{\pm 1.32}$	$29.49_{\pm 0.84}$	$38.84_{\pm 1.58}$	36.68
	DNA-GPT-M	41.16 \pm 1.63	$12.53_{\pm 2.06}$	$33.57_{\pm 2.15}$	29.09	47.10 \pm 0.93	$22.72_{\pm 1.50}$	$40.54_{\pm 1.98}$	36.79
	Likelihood	78.58 _{±0.41}	66.61 _{±0.99}	$71.42_{\pm0.49}$	72.2	$78.57_{\pm0.41}$	66.61 _{±0.99}	$71.42_{\pm 0.49}$	72.20
	Likelihood-M	78.15 ± 4.58	$67.06_{\pm 4.19}$	73.40±4.38	72.87	87.44±0.70	73.70 ± 0.57	81.58+0.93	80.91
	Log-Rank	$79.67_{\pm 0.46}$	$65.85_{\pm 0.94}$	$70.66_{\pm 0.40}$	72.06	$79.67_{\pm 0.46}$	65.85 ± 0.94	$70.66_{\pm 0.40}$	72.06
	Log-Rank-M	$90.23_{\pm 0.45}$	75.40 \pm 0.97	84.44 _{±0.52}	83.35	90.14±0.53	75.61 \pm 0.83	84.63±0.59	83.46
	Entropy	$66.21_{\pm 0.88}$	$63.09_{\pm 1.05}$	$60.73_{\pm 0.91}$	63.34	$66.21_{\pm 0.88}$	$63.09_{\pm 1.05}$	$60.73_{\pm 0.91}$	63.34
	Entropy-M	$68.61_{\pm 0.94}$	$66.83_{\pm 0.95}$	$67.93_{\pm 1.08}$	67.79	$68.91_{\pm 0.83}$	$66.86_{\pm 0.95}$	$67.94_{\pm 1.04}$	67.90
AUROC	DetectGPT	$49.21_{\pm 2.31}$	$49.64_{\pm 0.53}$	$50.72_{\pm 6.91}$	49.85	$47.63_{\pm 0.59}$	$49.78_{\pm 0.60}$	$56.81_{\pm 1.41}$	51.40
	DetectGPT-M	$73.50_{\pm 9.13}$	$62.93_{\pm 4.59}$	$61.30_{\pm 15.58}$	65.91	$77.77_{\pm 4.05}$	$60.34_{\pm 3.37}$	$76.40_{\pm 1.66}$	71.50
	FastGPT	$67.72_{\pm 1.02}$	$58.50_{\pm 1.22}$	$56.68_{\pm 1.21}$	60.97	$67.72_{\pm 1.02}$	$58.50_{\pm 1.22}$	56.68 + 1.21	60.97
	FastGPT-M	$71.60_{\pm 6.86}$	$61.26_{\pm 5.62}$	$66.05_{\pm 9.70}$	66.30	$62.35_{\pm 0.49}$	$51.87_{\pm 1.38}$	$68.88_{\pm 1.00}$	61.03
	DNA-GPT	72.53 + 1.19	$66.07_{\pm 1.38}$	68.35 ± 0.98	68.98	$79.00_{\pm 0.86}$	72.27 \pm 0.81	77.96±0.94	76.41
	DNA-GPT-M	$75.37_{\pm 1.13}$	66.19 $_{\pm 1.40}$	$71.77_{\pm 0.79}$	71.11	$80.09_{\pm 1.26}$	$71.09_{\pm 1.12}$	$77.15_{\pm 1.02}$	76.11

Table 10: Performance on Reuters dataset. The detection models are trained on text generated by GPT4All.

GP14AII.								
Metric	Method	GPT4All	ChatGPT	Dolly	ChatGLM	Claude	ChatGPT-turbo	Avg.
	Likelihood	$19.47_{\pm 2.45}$	$80.89_{\pm 2.17}$	$12.44_{\pm 1.38}$	$94.36_{\pm 1.15}$	$18.89_{\pm 3.41}$	$90.22_{\pm 1.23}$	52.71
	Likelihood-M	$61.42_{\pm 5.74}$	$88.44_{\pm 2.19}$	$35.56_{\pm 2.86}$	$95.24_{\pm 1.16}$	$32.36_{\pm 5.93}$	92.76 $_{\pm 0.71}$	67.63
	Log-Rank	$29.91_{\pm 3.27}$	$86.22_{\pm 1.87}$	$15.20_{\pm 1.70}$	$97.33_{\pm 0.64}$	$24.58_{\pm 3.35}$	$93.16_{\pm 1.07}$	57.73
	Log-Rank-M	$73.82_{\pm 3.32}$	89.56 \pm 1.79	$38.31_{\pm 2.26}$	$96.71_{\pm 1.62}$	$35.47_{\pm 4.63}$	$94.49_{\pm 0.62}$	71.39
	Entropy	$6.18_{\pm 1.40}$	$0.04_{\pm 0.09}$	$13.07_{\pm 1.76}$	$0.22_{\pm 0.20}$	$0.00_{\pm 0.00}$	$0.22_{\pm 0.00}$	3.29
TPR@FPR-1%	Entropy-M	$17.42_{\pm 2.84}$	$0.36_{\pm 0.30}$	12.36 + 2.86	$0.98_{\pm0.87}$	$0.00_{\pm 0.00}$	$0.40_{\pm 0.38}$	5.25
IFK@FFK-1%	DetectGPT	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.44_{\pm 0.40}$	$0.04_{\pm 0.09}$	0.08
	DetectGPT-M	$40.89_{\pm 6.07}$	$15.33_{\pm 8.81}$	$46.40_{\pm 2.46}$	$40.00_{\pm 8.10}$	$2.40_{\pm 0.38}$	$6.40_{\pm 2.50}$	25.24
	FastGPT	3.56 ± 3.08	$0.04_{\pm 0.09}$	12.18 ± 6.70	$0.09_{\pm 0.11}$	$0.09_{\pm 0.11}$	$0.04_{\pm 0.09}$	2.67
	FastGPT-M	$61.69_{\pm 7.18}$	$21.78_{\pm 6.77}$	$65.24_{\pm 5.12}$	$31.78_{\pm 9.97}$	$1.38_{\pm0.84}$	$1.51_{\pm 0.55}$	30.56
	DNA-GPT	$85.20_{\pm 6.79}$	$92.58_{\pm 3.94}$	$68.49_{\pm 4.37}$	$96.93_{\pm 3.14}$	$79.24_{\pm 4.93}$	$95.29_{\pm 3.53}$	86.29
	DNA-GPT-M	97.51 \pm 1.31	98.40 \pm 0.88	81.29 ± 3.75	99.87 \pm 0.11	89.07 \pm 3.57	99.24 \pm 0.30	94.23
	Likelihood	75.19 _{±1.39}	97.55 ± 0.33	58.77±1.76	99.62 _{±0.11}	$86.45_{\pm 0.70}$	98.42+0.25	86.00
	Likelihood-M	$93.74_{\pm 1.23}$	$98.42_{\pm 0.26}$	$77.05_{\pm 0.53}$	$99.61_{\pm 0.11}$	$90.68_{\pm 0.56}$	$98.78_{\pm 0.33}$	93.05
	Log-Rank	$78.39_{\pm 1.23}$	$97.77_{\pm 0.32}$	$58.22_{\pm 2.03}$	$99.61_{\pm 0.19}$	$85.68_{\pm0.67}$	$98.67_{\pm 0.24}$	86.39
	Log-Rank-M	$95.50_{\pm 0.81}$	$98.60_{\pm 0.23}$	$76.60_{\pm 1.33}$	99.64 \pm 0.29	$90.45_{\pm 0.69}$	$98.99_{\pm 0.31}$	93.30
	Entropy	$72.48_{\pm 1.24}$	$35.35_{\pm 1.65}$	$61.92_{\pm 1.32}$	$50.96_{\pm 1.57}$	$23.76_{\pm 1.05}$	$17.55_{\pm 0.92}$	43.67
AUROC	Entropy-M	$70.31_{\pm 1.98}$	$34.84_{\pm 2.95}$	$50.09_{\pm 1.57}$	$46.62_{\pm 1.12}$	$27.55_{\pm 2.60}$	$29.08_{\pm 4.18}$	43.08
AUROC	DetectGPT	$75.62_{\pm 1.20}$	$57.18_{\pm 1.91}$	$77.25_{\pm 0.66}$	$61.97_{\pm 0.90}^{-}$	$58.01_{\pm 1.75}$	$20.42_{\pm 1.07}$	58.41
	DetectGPT-M	$87.58_{\pm0.93}$	$78.19_{\pm 2.33}$	$86.29_{\pm 1.18}$	89.73 $_{\pm 0.97}$	$58.40_{\pm 1.54}$	$65.72_{\pm 1.36}$	77.65
	FastGPT	$76.50_{\pm 1.39}$	$39.56_{\pm 0.54}$	$82.45_{\pm0.47}$	$33.21_{\pm 1.82}$	$41.31_{\pm 0.89}$	$11.29_{\pm 0.49}$	47.39
	FastGPT-M	$96.46_{\pm0.31}$	90.21 $_{\pm 0.75}$	$96.23_{\pm 0.36}$	$92.36_{\pm 0.74}$	$66.98_{\pm 1.06}$	$56.28_{\pm 1.80}$	83.09
	DNA-GPT	$99.51_{\pm 0.14}$	$99.72_{\pm 0.07}$	$98.41_{\pm 0.21}$	$99.85_{\pm 0.05}$	$99.21_{\pm 0.14}$	$99.81_{\pm 0.06}$	99.42
	DNA-GPT-M	99.80 $_{\pm 0.09}$	99.88 ± 0.07	$98.39_{\pm 0.17}^{\pm 0.21}$	$99.94_{\pm 0.06}$	99.53 $_{\pm 0.09}$	$99.89_{\pm 0.09}$	99.57

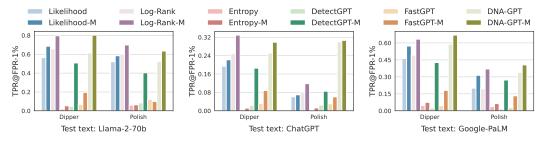


Figure 22: Detection performance concerning TPR@FPR-1% under different paraphrasing texts (Dipper and Polish). All detectors are trained on Llama-2-70b texts.

Table 11: Performance on Reuters dataset. The detection models are trained on text generated by ChatGPT

Metric	Method	GPT4All	ChatGPT	Dolly	ChatGLM	Claude	ChatGPT-turbo	Avg.
	Likelihood	19.78 ± 2.21	$81.07_{\pm 2.05}$	$12.53_{\pm 1.35}$	$94.36_{\pm 1.15}$	$19.42_{\pm 3.17}$	$90.31_{\pm 1.11}$	52.91
	Likelihood-M	$59.73_{\pm 8.56}$	$88.98_{\pm 1.58}$	$35.91_{\pm 3.82}$	$96.00_{\pm 2.14}$	$26.71_{\pm 5.93}$	$93.24_{\pm 0.94}$	66.76
	Log-Rank	$29.96_{\pm 3.29}$	$86.22_{\pm 1.87}$	$15.20_{\pm 1.70}$	$97.33_{\pm 0.64}$	$24.58_{\pm 3.35}$	$93.20_{\pm 1.13}$	57.75
	Log-Rank-M	$63.91_{\pm 8.04}$	$90.62_{\pm 2.38}$	$34.58_{\pm 4.71}$	$98.13_{\pm 0.46}$	$29.73_{\pm 8.71}$	$94.53_{\pm 0.67}$	68.59
	Entropy	$0.09_{\pm 0.11}$	$2.49_{\pm 1.25}$	2.04 ± 0.57	1.69 ± 0.54	4.36 + 1.81	9.56 ± 2.90	3.37
TPR@FPR-1%	Entropy-M	$7.07_{\pm 3.05}$	$17.47_{\pm 5.59}$	$13.60_{\pm 2.23}$	$9.02_{\pm 3.21}$	$18.09_{\pm 4.26}$	$33.29_{\pm 7.84}$	16.42
11 K@11 K-1 /0	DetectGPT	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.44_{\pm 0.40}$	$0.04_{\pm 0.09}$	0.08
	DetectGPT-M	$22.76_{\pm 25.66}$	$47.91_{\pm 28.49}$	11.69 $_{\pm 12.56}$	$23.64_{\pm 28.60}$	$8.27_{\pm 5.92}$	$71.24_{\pm 17.28}$	30.92
	FastGPT	$0.31_{\pm 0.23}$	2.04 ± 0.86	$0.04_{\pm 0.09}$	$3.16_{\pm 1.03}$	$1.42_{\pm 0.57}$	$23.51_{\pm 4.15}$	5.08
	FastGPT-M	$56.04_{\pm 11.84}$	$17.07_{\pm 10.24}$	$61.51_{\pm 9.34}$	$26.67_{\pm 13.24}$	0.89 ± 0.92	1.29 ± 0.68	27.24
	DNA-GPT	$29.82_{\pm 18.89}$	$69.78_{\pm 22.51}$	$15.64_{\pm 12.02}$	$83.47_{\pm 14.87}$	$21.96_{\pm 19.50}$	$82.22_{\pm 15.63}$	50.48
	DNA-GPT-M	$86.80_{\pm 5.37}$	$96.76_{\pm 1.16}$	59.33 _{±8.10}	99.29 $_{\pm 0.38}^{-}$	61.91 $_{\pm 13.48}$	97.51 $_{\pm 1.24}$	83.60
	Likelihood	$75.19_{\pm 1.39}$	97.55 ± 0.33	58.77±1.76	$99.62_{\pm0.11}$	86.46+0.70	$98.42_{\pm 0.25}$	86.00
	Likelihood-M	$93.97_{\pm 1.27}$	$98.68_{\pm 0.41}$	$79.54_{\pm 1.12}$	$99.70_{\pm 0.24}$	$90.39_{\pm 1.08}$	$98.94_{\pm 0.23}$	93.54
	Log-Rank	$78.39_{\pm 1.23}$	$97.77_{\pm 0.32}$	$58.22_{\pm 2.03}^{\pm 1.12}$	$99.61_{\pm 0.19}$	$85.68_{\pm 0.67}$	$98.67_{\pm 0.24}$	86.39
	Log-Rank-M	$93.41_{\pm 1.37}$	$98.69_{\pm 0.25}$	$76.76_{\pm 1.45}$	99.75 $_{\pm 0.20}$	$89.55_{\pm 1.04}$	$99.03_{\pm 0.27}$	92.86
	Entropy	$27.52_{\pm 1.24}$	$64.65_{\pm 1.65}$	$38.08_{\pm 1.32}$	$49.04_{\pm 1.57}$	$76.24_{\pm 1.05}$	$82.45_{\pm 0.92}$	56.33
AUROC	Entropy-M	$55.30_{\pm 2.65}$	$80.07_{\pm 1.89}$	$56.54_{\pm 2.20}$	69.75 ± 2.62	$83.33_{\pm 1.09}$	$90.46_{\pm 1.17}$	72.58
AUROC	DetectGPT	$75.62_{\pm 1.20}$	$57.18_{\pm 1.91}$	$77.25_{\pm 0.66}^{-}$	$61.97_{\pm 0.90}$	$58.01_{\pm 1.75}$	$20.42_{\pm 1.07}$	58.41
	DetectGPT-M	$84.60_{\pm 11.69}$	$92.37_{\pm 5.63}$	$69.47_{\pm 12.90}$	$80.47_{\pm 13.77}$	$75.78_{\pm 9.49}$	$97.20_{\pm 1.23}$	83.32
	FastGPT	$23.50_{\pm 1.39}$	$60.44_{\pm 0.54}$	$17.55_{\pm0.47}$	$66.79_{\pm 1.82}$	$58.69_{\pm 0.89}$	88.71 $_{\pm 0.49}$	52.61
	FastGPT-M	$96.25_{\pm 0.33}$	89.97 $_{\pm 0.74}$	$96.04_{\pm0.33}$	$92.12_{\pm 0.76}$	$66.72_{\pm 1.11}$	$56.04_{\pm 1.73}$	82.86
	DNA-GPT	$97.34_{\pm 0.54}$	$99.17_{\pm 0.30}$	95.82 ± 0.62	$99.45_{\pm0.23}$	$97.21_{\pm 0.52}$	99.38 ± 0.26	98.06
	DNA-GPT-M	99.42 $_{\pm 0.15}$	99.78 \pm 0.09	$97.31_{\pm 0.24}$	$99.89_{\pm0.10}$	$98.78_{\pm0.24}$	99.82 $_{\pm 0.12}$	99.16

Table 12: Performance on Reuters dataset. The detection models are trained on text generated by ChatGPT-turbo.

CnatGP1-tur	00.							
Metric	Method	GPT4All	ChatGPT	Dolly	ChatGLM	Claude	ChatGPT-turbo	Avg.
	Likelihood Likelihood-M	$19.82_{\pm 2.17}$ 66.62 _{± 5.21}	$81.07_{\pm 2.05}$ 90.13 _{\pm 1.15}	$12.53_{\pm 1.35}$ $38.62_{\pm 2.74}$	$94.44_{\pm 1.10}$ $95.96_{\pm 1.29}$	$19.42_{\pm 3.17}$ 29.60 _{±2.32}	$90.27_{\pm 1.07}$ $93.91_{\pm 1.03}$	52.93 69.14
	Log-Rank	$29.96_{\pm 3.29}$	$86.22_{\pm 1.87}$	$15.20_{\pm 1.70}$	$97.33_{\pm 0.64}$	$24.58_{\pm 3.35}$	$93.20_{\pm 1.13}$	57.75
	Log-Rank-M	60.18 ± 8.59	$90.27_{\pm 1.98}$	$33.07_{\pm 3.84}$	$98.22_{\pm 0.40}$	27.29 ± 6.31	$94.00_{\pm 0.40}$	67.17
	Entropy	$0.09_{\pm 0.11}$	$2.49_{\pm 1.25}$	$2.04_{\pm 0.57}$	$1.69_{\pm 0.54}$	$4.36_{\pm 1.81}$	$9.60_{\pm 2.94}$	3.38
TPR@FPR-1%	Entropy-M	$5.11_{\pm 2.70}$	$16.18_{\pm 4.91}$	$12.27_{\pm 1.51}$	$7.69_{\pm 2.70}$	17.42 \pm 3.93	$31.69_{\pm 6.61}$	15.06
IFK@IFK-1%	DetectGPT	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.62_{\pm 0.53}$	$9.24_{\pm 2.65}$	1.64
	DetectGPT-M	$6.22_{\pm 3.99}$	$26.71_{\pm 14.32}$	$2.49_{\pm 0.97}$	$4.04_{\pm 3.33}$	$4.18_{\pm 2.22}$	$60.67_{\pm 14.38}$	17.39
	FastGPT	$0.31_{\pm 0.23}$	$2.04_{\pm 0.86}$	$0.04_{\pm 0.09}$	$3.16_{\pm 1.03}$	$1.42_{\pm 0.57}$	$23.51_{\pm 4.15}$	5.08
	FastGPT-M	$0.00_{\pm 0.00}$	$0.22_{\pm 0.28}$	$0.04_{\pm 0.09}$	$0.53_{\pm 0.33}$	$0.62_{\pm 0.29}$	$7.38_{\pm 2.42}$	1.47
	DNA-GPT	2.09 ± 1.29	$4.31_{\pm 1.94}$	$1.69_{\pm 1.21}$	$5.51_{\pm 1.67}$	1.96 ± 1.60	$5.24_{\pm 2.29}$	3.47
	DNA-GPT-M	$5.82_{\pm 1.06}$	$6.44_{\pm 0.97}$	$3.56_{\pm 1.27}$	$6.49_{\pm 0.87}$	$3.78_{\pm 1.29}$	$6.67_{\pm 0.84}$	5.46
	Likelihood	$75.20_{\pm 1.39}$	$97.55_{\pm0.33}$	$58.77_{\pm 1.77}$	$99.62_{\pm0.11}$	$86.46_{\pm0.70}$	$98.42_{\pm0.25}$	86.00
	Likelihood-M	$94.79_{\pm 0.60}$	$98.72_{\pm 0.40}$	79.93 $_{\pm 1.44}$	99.72 \pm 0.16	$90.83_{\pm 0.53}$	$98.96_{\pm 0.26}$	93.82
	Log-Rank	$78.39_{\pm 1.23}$	$97.77_{\pm0.32}$	$58.22_{\pm 2.03}$	$99.61_{\pm 0.19}$	$85.68_{\pm0.67}$	$98.67_{\pm0.24}$	86.39
	Log-Rank-M	92.21 $_{\pm 2.45}$	$98.63_{\pm 0.23}$	76.27 $_{\pm 1.24}$	99.73 $_{\pm 0.21}$	$89.02_{\pm 0.89}$	$98.99_{\pm 0.27}$	92.48
	Entropy	$27.52_{\pm 1.24}$	$64.65_{\pm 1.65}$	$38.08_{\pm 1.32}$	$49.04_{\pm 1.57}$	$76.24_{\pm 1.05}$	$82.45_{\pm 0.92}$	56.33
AUROC	Entropy-M	$54.11_{\pm 3.57}$	79.38 $_{\pm 1.96}$	$55.94_{\pm 1.82}$	$68.88_{\pm 2.66}$	$83.14_{\pm 1.07}$	89.91 $_{\pm 0.91}$	71.89
AUROC	DetectGPT	24.38 ± 1.20	$42.82_{\pm 1.91}$	$22.75_{\pm 0.66}$	$38.03_{\pm 0.90}$	$41.99_{\pm 1.75}$	$79.58_{\pm 1.07}$	41.59
	DetectGPT-M	79.17 $_{\pm 7.47}$	89.36 $_{\pm 4.32}$	$59.50_{\pm 5.37}$	$71.41_{\pm 7.62}$	$73.92_{\pm 8.27}$	$96.53_{\pm 1.14}$	78.32
	FastGPT	$23.50_{\pm 1.39}$	$60.44_{\pm 0.54}$	$17.55_{\pm 0.47}$	66.79 ± 1.82	58.69 ±0.89	88.71 $_{\pm 0.49}$	52.61
	FastGPT-M	$12.73_{\pm 1.55}$	$35.93_{\pm 2.31}$	$9.43_{\pm 0.56}$	$36.46_{\pm 3.11}$	$50.18_{\pm 2.00}$	$76.28_{\pm 1.01}$	36.83
	DNA-GPT	$45.62_{\pm 1.73}$	$63.85_{\pm 1.00}$	$35.96_{\pm 1.28}$	$71.24_{\pm 1.14}$	$43.56_{\pm 1.49}$	$70.63_{\pm 1.33}$	55.14
	DNA-GPT-M	79.05 $_{\pm 1.13}$	$87.32_{\pm 0.51}$	$64.13_{\pm 1.27}$	$91.42_{\pm 0.64}$	67.14 $_{\pm 1.48}$	89.37 $_{\pm 0.79}$	79.74

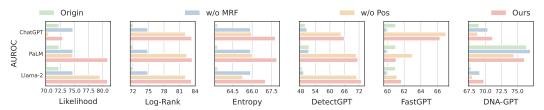


Figure 23: Ablation results on the DetectRL dataset. The y-axis represents the LLM text on which the detector was trained, and the x-axis represents the average performance across LLMs.

Table 13: Performance on Reuters dataset. The detection models are trained on text generated by ChatGLM

Metric	Method	GPT4All	ChatGPT	Dolly	ChatGLM	Claude	ChatGPT-turbo	Avg.
	Likelihood	$19.78_{\pm 2.13}$	81.07+2.05	$12.49_{\pm 1.33}$	$94.44_{\pm 1.10}$	$19.42_{\pm 3.17}$	$90.36_{\pm 1.06}$	52.93
	Likelihood-M	$41.42_{\pm 8.12}$	$86.84_{\pm 2.44}$	$31.47_{\pm 2.22}$	98.58 ± 0.23	20.84 ± 2.97	91.56 \pm 1.28	61.79
	Log-Rank	29.96 ± 3.29	86.22 ± 1.87	$15.20_{\pm 1.70}$	97.33 ± 0.64	24.58 ± 3.35	$93.20_{\pm 1.13}$	57.75
	Log-Rank-M	$62.40_{\pm 4.54}$	$90.49_{\pm 2.34}$	$33.82_{\pm 3.87}$	$98.22_{\pm 0.47}$	$30.44_{\pm 6.61}$	$94.13_{\pm 0.75}$	68.25
	Entropy	$4.89_{\pm 2.71}$	$0.36_{\pm 0.71}$	$10.49_{\pm 4.75}$	$0.44_{\pm 0.37}$	$0.58_{\pm 1.16}$	$1.38_{\pm 2.31}$	3.02
TPR@FPR-1%	Entropy-M	$6.13_{\pm 4.80}$	$17.20_{\pm 4.05}$	$13.11_{\pm 3.21}$	$9.73_{\pm 4.98}$	$18.76_{\pm 2.64}$	$33.87_{\pm 5.03}$	16.47
IFK@IFK-1%	DetectGPT	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.44_{\pm 0.40}$	$0.04_{\pm 0.09}$	0.08
	DetectGPT-M	$40.53_{\pm 6.43}$	14.18 ± 7.67	45.73 ± 2.50	$38.98_{\pm 8.03}$	2.18 ± 0.53	$5.47_{\pm 2.15}$	24.51
	FastGPT	$0.31_{\pm 0.23}$	$2.04_{\pm 0.86}$	$0.04_{\pm 0.09}$	$3.16_{\pm 1.03}$	$1.42_{\pm 0.57}$	$23.51_{\pm 4.15}$	5.08
	FastGPT-M	$56.04_{\pm 11.84}$	$17.07_{\pm 10.24}$	$61.51_{\pm 9.34}$	$26.67_{\pm 13.24}$	$0.89_{\pm 0.92}$	$1.29_{\pm 0.68}$	27.24
	DNA-GPT	$27.96_{\pm 18.31}$	$66.84_{\pm 18.61}$	$16.00_{\pm 12.45}$	$85.51_{\pm 12.79}$	$20.49_{\pm 17.96}$	$80.40_{\pm 13.38}$	49.53
	DNA-GPT-M	$82.89_{\pm 7.12}$	$94.93_{\pm 1.99}$	$56.31_{\pm 9.16}$	$98.98_{\pm0.18}$	$57.78_{\pm 15.40}$	$96.53_{\pm 1.23}$	81.24
	Likelihood	$75.19_{\pm 1.39}$	$97.55_{\pm0.33}$	58.77 _{+1.77}	99.62+0.11	86.46+0.70	98.42+0.25	86.00
	Likelihood-M	$87.15_{\pm 3.71}$	$98.44_{\pm 0.43}$	$77.92_{\pm 1.74}$	$99.95_{\pm 0.00}$	$87.99_{\pm 1.16}$	$98.69_{\pm 0.26}$	91.69
	Log-Rank	$78.39_{\pm 1.23}$	$97.77_{\pm 0.32}$	$58.22_{\pm 2.03}$	$99.61_{\pm 0.19}$	85.68 ± 0.67	$98.67_{\pm 0.24}$	86.39
	Log-Rank-M	$93.39_{\pm 0.75}$	$98.68_{\pm 0.24}$	76.11 \pm 1.61	99.75 \pm 0.19	89.70 $_{\pm 0.41}$	$98.99_{\pm 0.25}$	92.77
	Entropy	$62.76_{\pm 18.55}$	$40.53_{\pm 11.30}$	$56.16_{\pm 10.29}$	$49.39_{\pm 1.74}$	$33.61_{\pm 20.52}$	$30.23_{\pm 25.75}$	45.45
AUROC	Entropy-M	$54.30_{\pm 3.63}$	79.38 $_{\pm 1.33}$	$55.89_{\pm 1.46}$	$69.08_{\pm 2.00}$	$83.03_{\pm 0.44}$	89.77 $_{\pm 1.24}$	71.91
AUROC	DetectGPT	$75.62_{\pm 1.20}$	$57.18_{\pm 1.91}$	$77.25_{\pm 0.66}$	$61.97_{\pm 0.90}$	$58.01_{\pm 1.75}$	$20.42_{\pm 1.07}^{-}$	58.41
	DetectGPT-M	87.17 \pm 0.76	77.46 ± 1.85	86.14 ± 1.14	89.36 \pm 1.25	57.83 ± 1.55	$64.40_{\pm 1.89}$	77.06
	FastGPT	23.50 + 1.39	$60.44_{\pm 0.54}$	17.55 ± 0.47	$66.79_{\pm 1.82}$	58.69 ± 0.89	88.71 $_{\pm 0.49}$	52.61
	FastGPT-M	$96.27_{\pm 0.33}$	$89.98_{\pm 0.74}$	$96.07_{\pm0.33}$	$92.13_{\pm 0.76}$	66.74 ± 1.10	$56.05_{\pm 1.74}$	82.87
	DNA-GPT	$97.10_{\pm 0.48}$	$98.94_{\pm0.26}$	$95.56_{\pm0.57}$	$99.29_{\pm0.17}$	$96.95_{\pm0.46}$	$99.22_{\pm 0.22}$	97.85
	DNA-GPT-M	99.33 $_{\pm 0.15}$	99.72 $_{\pm 0.09}$	97.15 $_{\pm0.21}$	99.82 $_{\pm 0.10}$	$98.68_{\pm0.24}$	99.77 $_{\pm 0.12}$	99.08

Table 14: Performance on Reuters dataset. The detection models are trained on text generated by Dolly.

Metric	Method	GPT4All	ChatGPT	Dolly	ChatGLM	Claude	ChatGPT-turbo	Avg.
TPR@FPR-1%	Likelihood Likelihood-M Log-Rank Log-Rank-M Entropy Entropy-M DetectGPT DetectGPT-M FastGPT-M DNA-GPT-M DNA-GPT-M	19.78±2.21 59.96±5.45 6.62±11.36 66.76±7.01 6.18±1.40 2.40±0.74 0.00±0.00 42.93±25.42 3.56±3.08 61.69±7.18 78.71±8.56 93.24±5.57	$\begin{array}{c} 81.07 \pm 2.05 \\ \textbf{88.58} \pm 2.29 \\ \textbf{17.24} \pm 34.49 \\ \textbf{90.49} \pm 1.80 \\ \textbf{0.04} \pm 0.09 \\ \textbf{7.29} \pm 5.90 \\ \textbf{0.00} \pm 0.00 \\ \textbf{37.20} \pm 35.68 \\ \textbf{0.04} \pm 0.09 \\ \textbf{21.78} \pm 6.77 \\ \textbf{91.47} \pm 7.33 \\ \textbf{98.09} \pm 1.34 \end{array}$	$\begin{array}{c} 12.40\pm1.37\\ \textbf{35.42}\pm2.84\\ 10.49\pm2.42\\ \textbf{35.42}\pm4.26\\ \textbf{13.07}\pm1.76\\ 9.91\pm1.09\\ 0.00\pm0.00\\ \textbf{39.69}\pm19.14\\ 12.18\pm6.70\\ \textbf{65.24}\pm5.12\\ 70.00\pm10.40\\ \textbf{79.07}\pm7.29\\ \end{array}$	$\begin{array}{c} 94.44{\pm}1.10 \\ \textbf{95.91}{\pm}1.21 \\ 19.42{\pm}38.62 \\ \textbf{98.09}{\pm}0.72 \\ 0.22{\pm}0.20 \\ \textbf{3.16}{\pm}2.61 \\ 0.00{\pm}0.00 \\ \textbf{51.73}{\pm}37.94 \\ 0.09{\pm}0.11 \\ \textbf{31.78}{\pm}9.97 \\ \textbf{95.51}{\pm}4.74 \\ \textbf{99.87}{\pm}0.27 \end{array}$	19.38±3.23 31.96±6.22 4.93±9.87 31.78±6.22 0.00±0.00 9.24±7.67 0.44±0.40 6.27±7.20 0.09±0.11 1.38±0.84 76.76±8.45	90.27 ± 1.07 92.93 ± 0.77 18.58 ± 37.16 94.31 ± 0.39 0.22 ± 0.00 16.13 ± 13.02 0.04 ± 0.09 31.56 ± 34.65 0.04 ± 0.09 1.51 ± 0.55 94.80 ± 5.80 98.31 ± 1.35	52.89 67.46 12.88 69.47 3.29 8.02 0.08 34.90 2.67 30.56 84.54 92.30
	Likelihood	75.19 _{±1.39}	97.55±0.33	58.77±1.76	99.62±0.10	86.45±0.70	98.42+0.25	86.00
	Likelihood-M Log-Rank Log-Rank-M	$93.42_{\pm 0.91}$ $32.84_{\pm 22.66}$ $94.18_{\pm 0.71}$	$98.45_{\pm 0.29}$ $21.28_{\pm 38.17}$ $98.71_{\pm 0.19}$	$77.41_{\pm 1.04}$ $43.51_{\pm 5.44}$ $76.52_{\pm 1.64}$	$99.67_{\pm 0.16}$ $20.22_{\pm 39.67}$ $99.73_{\pm 0.23}$	$90.57_{\pm 0.61}$ $28.30_{\pm 28.33}$ $89.98_{\pm 0.70}$	$98.78_{\pm 0.33}$ $20.77_{\pm 38.91}$ $99.02_{\pm 0.31}$	93.05 27.82 93.02
	Entropy Entropy-M	72.48 \pm 1.24 49.80 \pm 2.95	$35.35_{\pm 1.65}$ $56.16_{\pm 26.44}$	$61.92_{\pm 1.32}$ $52.47_{\pm 1.86}$	$50.96_{\pm 1.57}$ $53.09_{\pm 16.63}$	23.76 ± 0.70 23.76 ± 1.05 56.78 ± 31.30	$17.55_{\pm 0.92}$ $58.33_{\pm 37.31}$	43.67 54.44
AUROC	DetectGPT DetectGPT-M	$75.62_{\pm 1.20}$ 80.34 _{+21.91}	$57.18_{\pm 1.91}$ $73.22_{\pm 28.90}$	$77.25_{\pm 0.66}$ 81.89 _{±11.41}	$61.97_{\pm 0.90}$ $91.04_{\pm 8.24}$	$58.01_{\pm 1.75}$ $57.53_{\pm 18.26}$	$20.42_{\pm 1.07}$ 66.30 _{± 30.95}	58.41 75.05
	FastGPT-M	$76.50_{\pm 1.39}$ 96.47 $_{\pm 0.30}$	$39.56_{\pm 0.54}$ $90.23_{\pm 0.73}$	$82.45_{\pm 0.47}$ 96.24 $_{\pm 0.37}$	$33.21_{\pm 1.82}$ $92.37_{\pm 0.72}$	$41.31_{\pm 0.89}$ 67.04 $_{\pm 1.07}$	$11.29_{\pm 0.49}$ 56.34 $_{\pm 1.90}$	47.39 83.12
	DNA-GPT DNA-GPT-M	$99.19_{\pm 0.21}$ $99.66_{\pm 0.19}$	$99.74_{\pm 0.09}$ $99.84_{\pm 0.13}$	$98.90_{\pm 0.24}$ $98.81_{\pm 0.22}$	$99.83_{\pm 0.07}$ $99.91_{\pm 0.10}$	$99.21_{\pm 0.16}$ $99.47_{\pm 0.20}$	$99.80_{\pm 0.08}$ $99.86_{\pm 0.14}$	99.45 99.59

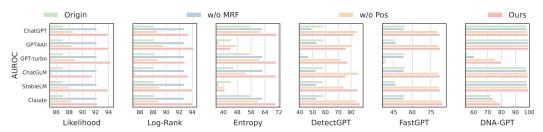


Figure 24: Ablation results on the Reuters dataset. The y-axis represents the LLM text on which the detector was trained, and the x-axis represents the average performance across LLMs.

Table 15: Performance on Reuters dataset. The detection models are trained on text generated by Claude.

claude.								
Metric	Method	GPT4All	ChatGPT	ChatGPT-turbo	ChatGLM	StableLM	Claude	Avg.
	Likelihood	19.82+2.17	$81.07_{\pm 2.05}$	$12.53_{\pm 1.35}$	94.44 ±1.10	$19.42_{\pm 3.17}$	$90.36_{\pm 1.06}$	52.94
	Likelihood-M	60.18 \pm 5.26	$84.84_{\pm 2.91}$	$32.13_{\pm 2.78}$	$91.78_{\pm 1.51}$	34.58 ± 4.58	91.47 $_{\pm 1.34}$	65.83
	Log-Rank	$29.96_{\pm 3.29}$	$86.22_{\pm 1.87}$	$15.20_{\pm 1.70}$	97.33 $_{\pm 0.64}$	24.58 + 3.35	$93.20_{\pm 1.13}$	57.75
	Log-Rank-M	$71.33_{\pm 4.00}$	$89.82_{\pm 2.47}$	$37.60_{\pm 1.64}$	$97.07_{\pm 1.44}$	$34.84_{\pm 5.56}$	$94.36_{\pm0.87}$	70.84
	Entropy	$0.09_{\pm 0.11}$	$2.49_{\pm 1.25}$	$2.04_{\pm 0.57}$	$1.69_{\pm 0.54}$	$4.36_{\pm 1.81}$	$9.60_{\pm 2.94}$	3.38
TPR@FPR-1%	Entropy-M	$5.51_{\pm 2.56}$	17.42 ± 5.92	$13.24_{\pm 2.62}$	$7.91_{\pm 2.58}$	$18.00_{\pm 4.28}$	$34.09_{\pm 8.23}$	16.03
IPK@FPK-1%	DetectGPT	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.44_{\pm 0.40}$	$0.04_{\pm 0.09}$	0.08
	DetectGPT-M	$39.91_{\pm 20.48}$	$51.60_{\pm 22.14}$	$28.80_{\pm 20.74}$	$57.47_{\pm 35.18}$	$6.00_{\pm 3.22}$	$78.53_{\pm 8.07}$	43.72
	FastGPT	$0.31_{\pm 0.23}$	2.04 ± 0.86	$0.04_{\pm 0.09}$	$3.16_{\pm 1.03}$	1.42 ± 0.57	$23.51_{\pm 4.15}$	5.08
	FastGPT-M	59.96 _{+10.36}	$33.11_{\pm 28.98}$	$56.09_{\pm 16.45}$	$42.44_{\pm 27.14}$	$8.62_{+15.04}$	$19.64_{\pm 36.52}$	36.64
	DNA-GPT	$3.64_{\pm 1.52}$	$6.58_{\pm 1.66}$	$2.18_{\pm 1.15}$	$7.78_{\pm 0.95}^{-}$	$2.44_{\pm 2.03}$	$7.20_{\pm 1.15}$	4.97
	DNA-GPT-M	6.67 _{±2.96}	7.16 $_{\pm 3.18}$	$4.22_{\pm 2.00}^{-}$	$7.51_{\pm 3.26}$	$5.38_{\pm 2.87}$	$7.16_{\pm 3.17}$	6.35
	Likelihood	75.19+1.39	$97.55_{\pm0.33}$	58.77 _{+1.77}	99.62 _{±0.10}	86.45+0.70	98.42 _{+0.25}	86.00
	Likelihood-M	91.12+2.58	97.69 _{±0.49}	$69.45_{\pm 3.64}$	$98.38_{\pm 0.57}$	$90.57_{\pm 0.59}$	$98.54_{\pm 0.43}$	90.96
	Log-Rank	$78.40_{\pm 1.23}$	$97.77_{\pm0.32}$	$58.22_{\pm 2.03}$	$99.61_{\pm 0.19}$	$85.68_{\pm 0.67}$	$98.67_{\pm 0.24}$	86.39
	Log-Rank-M	$95.06_{\pm 0.64}$	$98.59_{\pm 0.28}$	$76.35_{\pm 1.49}$	99.65 \pm 0.22	$90.36_{\pm 0.68}$	98.98 ± 0.27	93.16
	Entropy	$27.52_{\pm 1.24}$	$64.65_{\pm 1.65}$	$38.08_{\pm 1.32}$	$49.04_{\pm 1.57}$	$76.24_{\pm 1.05}$	$82.45_{\pm 0.92}$	56.33
AUROC	Entropy-M	$53.47_{\pm 2.35}$	79.33 $_{\pm 1.73}$	$55.92_{\pm 1.77}$	$68.50_{\pm 2.29}$	$83.03_{\pm 1.03}$	$90.09_{\pm 0.94}$	71.72
AUROC	DetectGPT	$75.62_{\pm 1.20}$	$57.18_{\pm 1.91}$	$77.25_{\pm 0.66}$	$61.97_{\pm 0.90}$	$58.01_{\pm 1.75}$	$20.42_{\pm 1.07}$	58.41
	DetectGPT-M	91.23 _{±5.89}	$94.13_{\pm 1.77}$	$81.29_{\pm 8.98}$	$94.63_{\pm 5.95}$	$70.99_{\pm 6.72}$	$96.97_{\pm 1.09}$	88.21
	FastGPT	$23.50_{\pm 1.39}$	$60.44_{\pm 0.54}$	$17.55_{\pm0.47}$	$66.79_{\pm 1.82}$	$58.69_{\pm 0.89}$	88.71 $_{\pm 0.49}$	52.61
	FastGPT-M	$96.53_{\pm 0.51}$	91.64 $_{\pm 3.61}$	$94.46_{\pm 2.87}$	$93.63_{\pm 3.03}$	71.74 ± 10.00	$63.82_{\pm 17.64}$	85.30
	DNA-GPT	$65.77_{\pm 1.29}$	$78.38_{\pm 0.64}$	$54.85_{\pm 1.20}$	$82.40_{\pm 0.73}$	$65.30_{\pm 0.67}$	$81.64_{\pm 0.76}$	71.39
	DNA-GPT-M	78.84 ± 12.13	$83.32_{\pm 13.75}$	$68.13_{\pm 8.88}$	$85.53_{\pm 14.38}$	$72.95_{\pm 9.83}$	$84.29_{\pm 14.01}$	78.84

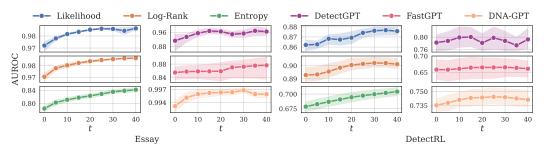


Figure 25: Detection performance of AUROC under different transition center t_0 .

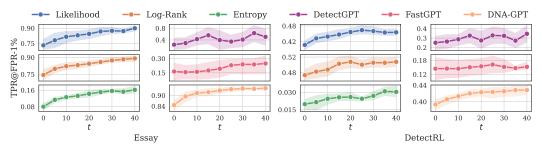


Figure 26: Detection performance of TPR@FPR-1% under different transition center t_0 .

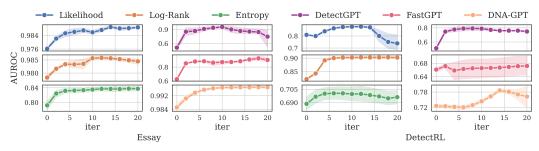


Figure 27: Detection performance of AUROC at different numbers of MRF layer iterations.

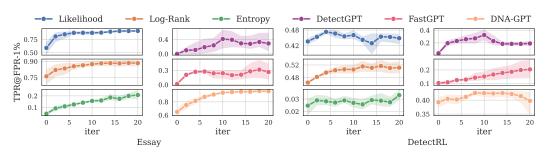


Figure 28: Detection performance of TPR@FPR-1% at different numbers of MRF layer iterations.

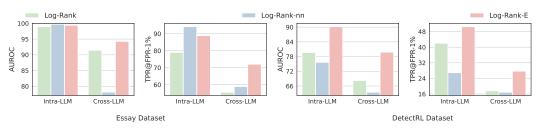


Figure 29: Comparison with NN-based methods. The detector used is Log-Rank.

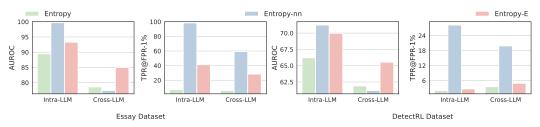


Figure 30: Comparison with NN-based methods. The detector used is Entropy.

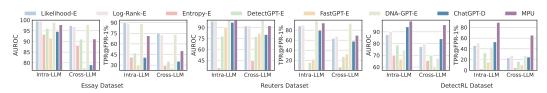


Figure 31: Comparison with Model-based detectors. Detectors are trained on GPT4All (Essay and Reuters) and Llama-2-70b (DetectRL).

Table 16.	Running t	time (s) o	f training and	inference phases.
Table 10.	Kummin	Time (S) O	u manning and	interence phases.

Method		Train			Inference	
	Essay	Reuters	DetectRL	Essay	Reuters	DetectRL
Likelihood	13.58	12.77	12.40	61.21	61.05	59.34
Likelihood-M	15.01	14.19	13.10	72.67	72.38	61.66
Log-Rank	15.98	14.38	14.61	75.60	72.99	66.20
Log-Rank-M	17.40	15.81	16.29	87.55	72.06	82.21
Entropy	13.26	12.93	13.79	64.49	63.69	63.19
Entropy-M	14.69	14.38	15.61	73.66	72.80	72.63
DetectGPT	204.32	216.36	370.44	924.63	982.55	1672.85
DetectGPT-M	206.44	218.73	373.48	929.46	984.70	1682.78
FastGPT	60.21	58.67	52.97	279.50	272.18	240.20
FastGPT-M	63.00	61.47	55.76	286.44	284.75	254.85
DNA-GPT	469.10	532.02	267.84	2113.65	2398.61	1207.18
DNA-GPT-M	471.42	534.47	271.11	2125.17	2405.37	1221.74