

MASH: Evading Black-Box AI-Generated Text Detectors via Style Humanization

Yongtong Gu¹, Songze Li^{1*}, Xia Hu²

¹Southeast University

²Shanghai Artificial Intelligence Laboratory

{yongtonggu, songzeli}@seu.edu.cn, huxia@pjlab.org.cn

Abstract

The increasing misuse of AI-generated texts (AIGT) has motivated the rapid development of AIGT detection methods. However, the reliability of these detectors remains fragile against adversarial evasions. Existing attack strategies often rely on white-box assumptions or demand prohibitively high computational and interaction costs, rendering them ineffective under practical black-box scenarios. In this paper, we propose Multi-stage Alignment for Style Humanization (MASH), a novel framework that evades black-box detectors based on style transfer. MASH sequentially employs style-injection supervised fine-tuning, direct preference optimization, and inference-time refinement to shape the distributions of AI-generated texts to resemble those of human-written texts. Experiments across 6 datasets and 5 detectors demonstrate the superior performance of MASH over 11 baseline evaders. Specifically, MASH achieves an average Attack Success Rate (ASR) of 92%, surpassing the strongest baselines by an average of 24%, while maintaining superior linguistic quality. Our code and data are publicly available at <https://github.com/githigher/MASH>.

1 Introduction

Large Language Models (LLMs), such as ChatGPT, Gemini, and Claude, have demonstrated remarkable capabilities across various domains, including social media (Macko et al., 2025), creative writing (Baek et al., 2025), and code generation (Li et al., 2025). However, the rapid proliferation of these capabilities has raised serious concerns across academia and industry over issues such as automated fake news, web content pollution, and academic plagiarism (Wang et al., 2024a; Zheng et al., 2025). To maintain the authenticity of the information ecosystem, detecting AI-generated text (AIGT) has become a critical issue. Recent studies

*Corresponding author.

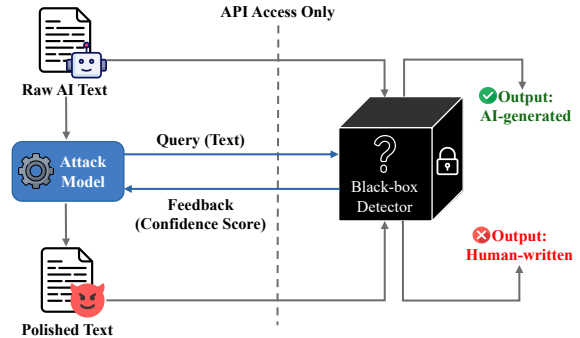


Figure 1: Illustration of an evasion attack against AIGT detectors in a black-box setting.

(Cheng et al., 2025; Russell et al., 2025) indicate that existing AIGT detectors have achieved remarkable accuracy on standard benchmarks.

Nevertheless, the robustness of these detectors remains fragile when facing malicious attacks. As illustrated in Figure 1, an attacker can rewrite the original generation to alter its stylistic pattern, thereby fooling the detector into misclassifying it as human-written. However, existing attack strategies often rely on white-box assumptions, leading to limited Attack Success Rates (ASR) in black-box settings. Furthermore, these methods face significant obstacles in real-world deployment, where detectors typically operate as black boxes with strict access limits. Major constraints include: (1) excessive query interactions during inference (Abad-Rocamora et al., 2024; Zhou et al., 2024); (2) high computational burdens (Fang et al., 2025); and (3) unrealistic requirements for detector gradients (Meng et al., 2025).

To address these challenges, we propose Multi-stage Alignment for Style Humanization (MASH), a novel black-box style transfer framework for detector evasion. Current detectors identify AI-generated text by leveraging inherent distribution shifts in semantic (Huang et al., 2024b) and statistical features (Hans et al., 2024). Consequently, we

posit that effective evasion necessitates humanizing AI-specific features through style transfer. MASH trains a paraphraser via a sequential pipeline: First, it employs Style-injection Supervised Fine-Tuning (Style-SFT) on open-source corpora, where the model learns by explicitly extracting style vectors from AI-generated and human-written text, respectively. Building on this foundation, the paraphraser is further fine-tuned using Direct Preference Optimization (DPO) (Rafailov et al., 2023) to adapt to the detector’s decision boundaries, thereby significantly boosting the ASR. Finally, adversarial refinement is applied during inference to guarantee the linguistic quality of the output.

In summary, our contributions are as follows:

- We propose MASH, a multi-stage black-box attack framework that sequentially integrates Style-SFT, DPO alignment, and inference-time refinement to achieve effective evasion.
- We demonstrate that a 0.1B parameter model, optimized via MASH, can surpass much larger models in evasion performance while requiring only limited interactions.
- Systematic evaluations across six domains confirm that MASH significantly outperforms state-of-the-art baselines, achieving superior attack success rates against both open-source and commercial detectors while maintaining superior linguistic quality.

2 Related Work

2.1 AI-generated Text Detection

Training-based Detectors. These methods typically formulate detection as a binary classification problem. Recent studies have incorporated adversarial training (Mao et al., 2024), Longformer architectures (Li et al., 2024), contrastive learning (Guo et al., 2024), and denoising reconstruction mechanisms (Huang et al., 2024b). Furthermore, such supervised approaches typically require extensive manual annotations.

Zero-shot Detectors. These methods eliminate the need for training by leveraging inherent statistical features. Representative approaches include GLTR (Gehrmann et al., 2019), which utilizes token ranking, and DetectGPT (Mitchell et al., 2023), which examines probability curve curvature. Other works include Fast-DetectGPT (Bao et al., 2023)

and Binoculars (Hans et al., 2024); the latter identifies machine text through ratio relationships between perplexity and cross-perplexity. This study focuses only on training-based and zero-shot detectors due to their practical prevalence in deployment.

2.2 Detection Evasion Methods

Perturbation-based Attacks. These methods disrupt detectors through character or word-level modifications. Early works like DeepWordBug (Gao et al., 2018) and TextFooler (Jin et al., 2020) utilize spelling errors or synonym substitutions. Recently, Charmer (Abad Rocamora et al., 2024) revitalizes this direction by optimizing character-level noise distributions. However, their ASR remains suboptimal when the budget for detector interactions is strictly constrained.

Prompt-based Attacks. These methods use contextual prompts to guide LLMs in generating evasion-prone text. SICO (Lu et al., 2024) employs optimization to find detection-minimizing prompts, while PromptAttack (Xu et al., 2024) uses prompts to mimic human styles. Despite producing high-quality text, their efficacy depends on the model’s instruction-following consistency and can be unstable against robust classifiers.

Paraphrase-based Attacks. These methods utilize specialized models to rephrase machine-generated text. DIPPER (Krishna et al., 2023) introduces a T5-XXL paraphraser that enables control over lexical and syntactic parameters. Subsequently, Sadasivan et al. (2023) proposed using recursive paraphrasing. Recent works leverage DPO for evasion, using detector confidence as a reward (Nicks et al., 2023; Wang et al., 2025). Building on this, Pedrotti et al. (2025) fine-tune models on diverse parallel data. However, these methods often necessitate white-box access to the source generator. In contrast, MASH is a black-box paradigm that humanizes text from any LLM without internal access to either the source generator or the target detector.

2.3 Text Style Transfer

Text style transfer (TST) aims to modify linguistic attributes while strictly preserving semantic fidelity, such as sentiment or formality, via learning disentangled latent representations with designated semantics (Hu et al., 2017; Shen et al., 2017). Recently, LLMs have shifted this paradigm toward in-context learning and parameter-efficient fine-tuning (Reif et al., 2022). Notably, recent studies reveal the intrinsic stylistic divergence between

AI-generated and human-written text (Jiang et al., 2025; Liu and Demberg, 2023). In this work, we extend TST to stylistic humanization, redefining detector evasion as a specialized “machine-to-human” style transfer task.

3 Methodology

3.1 Problem Formulation

The task of AI-generated text detection is formulated as a binary classification problem. Given a sequence $\mathbf{x} = (x_1, \dots, x_T)$ generated by an LLM via $P_{LLM}(\mathbf{x}) = \prod_{t=1}^T P(x_t|x_{<t})$, a detector D assigns a score $D(\mathbf{x}) \in [0, 1]$ representing the probability of AI. The final decision \hat{y} is:

$$\hat{y} = \mathbb{1}(D(\mathbf{x}) > \tau), \quad (1)$$

where $\mathbb{1}(\cdot)$ is the indicator function, τ is the threshold, and $\hat{y} \in \{1, 0\}$ denotes AI-generated and human-written text, respectively.

Threat Model. We consider a black-box setting where the attacker has no knowledge of the detector’s architecture, parameters, or gradients. Access is restricted to an Oracle providing scores $D(\mathbf{x})$ or labels \hat{y} . The objective is to learn a generator G_ϕ that transforms source text \mathbf{x}_{ai} into an adversarial example \mathbf{x}_{adv} to evade detection:

$$\begin{aligned} & \text{minimize} \quad \mathbb{E}_{\mathbf{x}_{ai} \sim P_{LLM}} [D(G_\phi(\mathbf{x}_{ai}))], \\ & \text{subject to} \quad \mathcal{S}(\mathbf{x}_{adv}, \mathbf{x}_{ai}) \geq \epsilon, \quad \mathcal{Q}(\mathbf{x}_{adv}) \geq \delta, \end{aligned} \quad (2)$$

where \mathcal{S} and \mathcal{Q} denote semantic consistency and fluency with thresholds ϵ and δ , respectively.

To achieve the adversarial objective in a black-box setting, we develop the Multi-stage Alignment for Style Humanization (MASH) framework. Our key idea is to reformulate detector evasion as a style transfer task. As illustrated in Figure 2, MASH bridges the stylistic gap between AI and human texts through four stages. We first synthesize parallel data via Inverse Data Construction, enabling Style-SFT to initialize the paraphraser with human writing patterns. Subsequently, DPO Alignment optimizes the paraphraser against detector boundaries, culminating in Inference-Time Adversarial Refinement to ensure linguistic quality.

3.2 MASH Methodology

Stage 1: Data Construction. Text style transfer is often limited by the scarcity of parallel training data. Our key insight is that while converting AI-generated text to human style is challenging, the

reverse direction is easy: LLMs can readily rewrite human-written text into machine style. Moreover, human-written text is abundantly available from open-source datasets. First, we collect raw texts from open-source datasets and retain only high-confidence human-written samples \mathbf{x}_{human} (where $D(\mathbf{x}_{human}) < \tau$). Next, for each human sample \mathbf{x}_{human} , we use an LLM to generate a semantically equivalent paraphrase \mathbf{x}_{ai} . These are further filtered to ensure strong machine style ($D(\mathbf{x}_{ai}) > \tau$). Consequently, we obtain a dataset consisting of N AI-human pairs $\mathcal{D}_{pair} = \{(\mathbf{x}_{ai}^{(i)}, \mathbf{x}_{human}^{(i)})\}_{i=1}^N$.

Stage 2: Style-injection Supervised Fine-Tuning.

Direct fine-tuning often leads to overfitting and semantic loss. To achieve controllable humanization, we propose a style-injection architecture based on pre-trained BART. We decouple style by introducing trainable embeddings $\mathbf{s}_{ai}, \mathbf{s}_{human} \in \mathbb{R}^d$. Given input \mathbf{x}_{ai} , the encoder produces content representation $\mathbf{H}_{content}$. A fusion layer then injects the selected style \mathbf{s}_{style} via linear projection:

$$\mathbf{H}_{fused}^{(t)} = \mathbf{W}_p \cdot [\mathbf{h}_{content}^{(t)}; \mathbf{s}_{style}] + \mathbf{b}_p, \quad (3)$$

where $[\cdot; \cdot]$ denotes the concatenation, and \mathbf{W}_p and \mathbf{b}_p represent the trainable weight matrix and bias vector of the projection layer, respectively.

The decoder generates the target sequence conditioned on the fused representation \mathbf{H}_{fused} . To balance style injection with semantic preservation, we employ a multi-task objective:

$$\mathcal{L}_{recon} = - \sum \log P_\theta(\mathbf{x}_{ai} | \mathbf{x}_{ai}, \mathbf{s}_{ai}), \quad (4)$$

$$\mathcal{L}_{trans} = - \sum \log P_\theta(\mathbf{x}_{human} | \mathbf{x}_{ai}, \mathbf{s}_{human}), \quad (5)$$

where \mathcal{L}_{recon} preserves semantic consistency by reconstructing the original AI-generated text, while \mathcal{L}_{trans} imparts human stylistic patterns by generating the target \mathbf{x}_{human} conditioned on \mathbf{s}_{human} . We minimize $\mathcal{L}_{SFT} = \lambda \mathcal{L}_{recon} + (1 - \lambda) \mathcal{L}_{trans}$ to inject human style while preserving semantics, ensuring robust initialization for alignment.

Stage 3: DPO Alignment. In the context of style transfer, SFT learns what human style looks like, but not how far the output has moved from AI style. To bridge this gap, we employ DPO, which uses the detector’s confidence as an implicit reward signal to directly optimize for crossing the style boundary. Formally, we posit that the human preference follows a Bradley-Terry model (Bradley and Terry, 1952) driven by an implicit reward function $r(x, y)$.

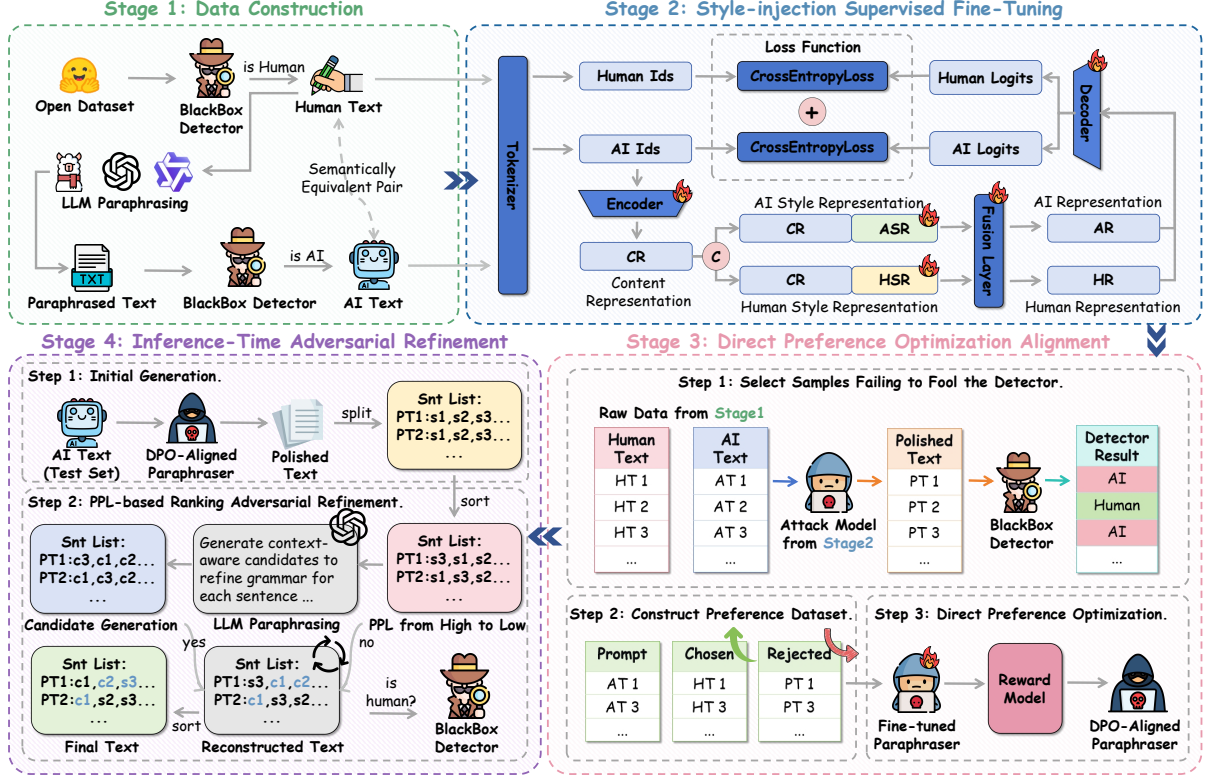


Figure 2: Overview of the proposed MASH framework. It comprises four stages: (1) Data Construction to synthesize parallel data; (2) Style-Injection SFT for supervised initialization; (3) DPO Alignment to optimize against detector boundaries; and (4) Inference-Time Adversarial Refinement to guarantee final text quality.

To align the optimization objective with our evasion goal, we define the reward as a scaled inversion of the detector’s confidence score $D(y)$:

$$r(x, y) = -C \cdot D(y), \quad C > 0. \quad (6)$$

As demonstrated in Appendix A.1, maximizing this reward forces the optimal policy π^* to converge towards regions where $D(y) \rightarrow 0$.

Hard Negative Mining. We construct the preference dataset by repurposing the parallel data \mathcal{D}_{pair} from Stage 1, eliminating the need for external data. For each source input x_{ai} , the ground-truth human text serves as the chosen response y_w . The rejected response y_l is sampled from the Stage 2 model $\pi_{SFT}(\cdot | x_{ai}, s_{human})$ and retained only if it fails to evade the detector (i.e., $D(y_l) > \tau$). As justified in Appendix A.2, this selection criterion maximizes the probabilistic gap between pairs, preventing vanishing gradients and ensuring robust optimization. The resulting dataset is formulated as:

$$\mathcal{D}_{DPO} = \left\{ (x^{(i)}, y_w^{(i)}, y_l^{(i)}) \mid D(y_l^{(i)}) > \tau \right\}_{i=1}^N. \quad (7)$$

Fine-Tuning with DPO. The theoretical justification for using Style-SFT to ensure support cover-

age is detailed in Appendix A.3. In practice, we bypass explicit reward modeling and directly optimize the policy. Following Nicks et al. (2023), we define the implicit log-ratio $h_\theta(y|x) = \beta \log \frac{\pi_\theta(y|x)}{\pi_{ref}(y|x)}$, yielding the adversarial loss:

$$\mathcal{L}_{DPO} = -\mathbb{E}_{\mathcal{D}_{DPO}} [\log \sigma (h_\theta(y_w|x) - h_\theta(y_l|x))]. \quad (8)$$

Minimizing this loss functions as implicit adversarial training, pushing the generation distribution from the AI region toward the human region based on the detector’s judgments.

Stage 4: Inference-Time Adversarial Refinement. We introduce a refinement module that enhances text quality while strictly preserving the detector’s human-written prediction.

Candidate Generation. We decompose the DPO-aligned text \hat{T} into a sentence sequence $S = \{s_1, \dots, s_n\}$. For each sentence s_i , an LLM polisher \mathcal{G} generates a candidate c_i conditioned on the original machine-generated content C_{ref} and a polishing instruction I , formulated as $c_i = \mathcal{G}(s_i | C_{ref}, I)$.

PPL-based Ranking Adversarial Refinement. To optimize query efficiency, we rank sentences by perplexity in descending order, prioritizing low-

fluency sentences. We apply a greedy strategy where s_i is replaced by c_i only if the detector D maintains a human prediction. This enhances fluency without compromising attack success.

4 Experiments

4.1 Experimental Setup

Datasets and Detectors. To ensure a comprehensive evaluation (Zheng et al., 2025), we evaluate MASH on six domains from two benchmarks: MGTBench (Essay, Reuters, WP) (He et al., 2024) and MGT-Academic (STEM, Social, Humanity) (Liu et al., 2025). We test against three open-source detectors (RoBERTa, Binoculars, SCRN) and two commercial APIs (Writer, Scribbr).

Baselines and Metrics. We benchmark against 11 state-of-the-art methods: five perturbation-based attacks (DeepWordBug (Gao et al., 2018), TextBugger (Li et al., 2018), TextFooler (Jin et al., 2020), Charmer (Abad Rocamora et al., 2024), HMGC (Zhou et al., 2024)), one prompt-based attack (PromptAttack (Xu et al., 2024)), and five paraphrase-based attacks (DIPPER (Krishna et al., 2023), DPO-Evader (Nicks et al., 2023), ToBlend (Huang et al., 2024a), CoPA (Fang et al., 2025), GradEscape (Meng et al., 2025)).

Following Meng et al. (2025); Wang et al. (2025), we evaluate evasion performance via Attack Success Rate (ASR) and text quality via Perplexity (PPL), BERTScore (Zhang et al., 2019), and GRUEN (Zhu and Bhat, 2020).

Implementation Details. MASH is initialized with BART-base (Lewis et al., 2020) and optimized via AdamW on a single NVIDIA RTX 3090 GPU. Detailed experimental configurations are provided in Appendix B.

4.2 Performance Evaluation

Attack Effectiveness. As shown in Table 1, MASH outperforms the strongest baselines by an average of 29.7% against the fine-tuned RoBERTa detector, while maintaining superior text quality. The limitations of perturbation-based attacks are illustrated in Figure 11 of Appendix C. Paraphrase-based methods also perform poorly. Among existing paraphrase-based baselines, DIPPER struggles with out-of-distribution (OOD) generalization, while CoPA’s indirect proxy-guidance fails to capture target-specific decision boundaries.

We further evaluate MASH against zero-shot detectors. As detailed in Table 2, MASH surpasses

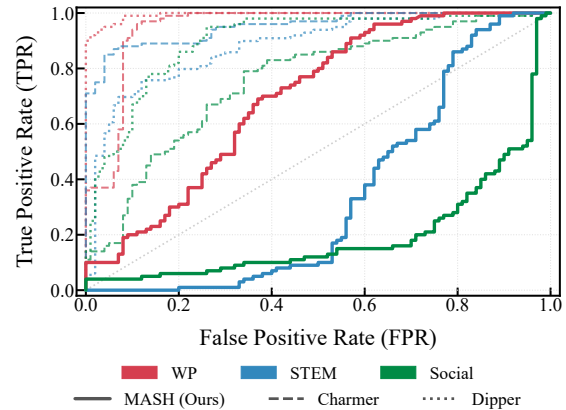


Figure 3: ROC curves comparing the evasion effectiveness of MASH against baseline attacks on the RoBERTa detector across the WP, STEM, and Social domains. Note that a lower TPR corresponds to a higher ASR.

the strongest baselines by an average of 19.0%. This suggests that transferring to human style naturally corrects the statistical patterns that zero-shot detectors rely on.

For commercial detectors, we assess the detection capabilities of Writer and Scribbr, which exhibit high precision across all six domains (see Table 6 in Appendix C), confirming their robustness. As shown in Table 3, MASH evades these detectors with an average ASR of 89%, while preserving semantic consistency. These results confirm that MASH generalizes effectively to real-world commercial systems under black-box settings.

Interpretation of AUROC. In real-world detection scenarios, the utility of a detector hinges on its capacity to maintain a high True Positive Rate (TPR) while rigorously suppressing the False Positive Rate (FPR) to prevent wrongful accusations against human authors. To evaluate this, we benchmarked our method against two highly efficient baselines, Charmer and DIPPER. Figure 3 illustrates the ROC curves under these three distinct attack modalities. Quantitatively, at a stringent 1% FPR threshold, MASH suppresses the TPR to near-zero levels, particularly in the Social and STEM domains. This confirms that MASH effectively neutralizes the detector’s discriminative capability, even under conservative operational constraints.

Efficiency Analysis. To evaluate practical feasibility, we exclude the Stage 4 refinement module (optional for users prioritizing ASR) and compare inference overhead. As reported in Table 4, perturbation-based attacks (e.g., Charmer) incur prohibitive query costs, while paraphrase-based

Method	MGTBench-Essay				MGTBench-Reuters				MGTBench-WP			
	ASR \uparrow	PPL \downarrow	GRUEN \uparrow	BS \uparrow	ASR \uparrow	PPL \downarrow	GRUEN \uparrow	BS \uparrow	ASR \uparrow	PPL \downarrow	GRUEN \uparrow	BS \uparrow
DeepWordBug	0.13	63.02	0.4703	0.9339	0.02	88.83	0.5240	0.9325	0.51	32.85	0.5238	0.9497
TextBugger	0.47	40.67	0.4572	0.9188	0.17	84.02	0.4043	0.8929	0.77	34.24	0.4936	0.9318
TextFooler	0.38	77.18	0.4465	0.9399	0.29	181.46	0.3018	0.8700	0.59	50.20	0.4249	0.9379
DIPPER	0.07	44.78	0.6481	0.9034	0.00	11.07	0.7180	0.9096	0.04	14.11	0.5773	0.9006
Toblend	0.00	7.03	0.3777	0.7525	0.00	6.88	0.5581	0.7242	0.19	6.59	0.4814	0.7312
PromptAttack	0.00	11.21	0.7776	0.9155	0.00	10.55	0.7909	0.9269	0.00	10.55	0.7256	0.9246
Charmer	0.45	38.00	0.5532	0.9664	0.05	114.57	0.5828	0.9624	0.62	17.26	0.5785	0.9759
DPO-Evader	0.00	5.14	0.1539	0.9343	0.00	4.53	0.1492	0.9257	0.00	5.41	0.2037	0.9230
HMGC	0.04	57.80	0.5883	0.9226	0.01	82.60	0.5469	0.9006	0.15	41.82	0.5763	0.9319
GradEscape	0.22	17.72	0.1345	0.9038	0.02	74.96	0.4811	0.8141	0.00	13.49	0.6516	0.9865
CoPA	0.01	29.42	0.6891	0.8750	0.00	41.64	0.6942	0.8769	0.17	25.47	0.6325	0.8807
Ours	0.95	18.90	0.6614	0.9004	0.73	9.09	0.6790	0.9015	0.90	20.80	0.6471	0.8974

Method	MGT-Academic-Humanity				MGT-Academic-Social Science				MGT-Academic-STEM			
	ASR \uparrow	PPL \downarrow	GRUEN \uparrow	BS \uparrow	ASR \uparrow	PPL \downarrow	GRUEN \uparrow	BS \uparrow	ASR \uparrow	PPL \downarrow	GRUEN \uparrow	BS \uparrow
DeepWordBug	0.07	30.09	0.4867	0.9520	0.11	32.59	0.5411	0.9514	0.07	22.64	0.4662	0.9543
TextBugger	0.15	40.51	0.4031	0.9319	0.11	56.87	0.4425	0.9364	0.06	39.93	0.4041	0.9333
TextFooler	0.11	72.33	0.3495	0.9289	0.10	136.89	0.3315	0.9121	0.07	81.24	0.3189	0.9255
DIPPER	0.57	14.67	0.5782	0.9031	0.51	13.88	0.6833	0.9061	0.55	12.47	0.6039	0.8789
PromptAttack	0.00	12.46	0.6893	0.9277	0.00	11.61	0.6904	0.9295	0.01	10.39	0.6529	0.9340
Charmer	0.73	14.54	0.5504	0.9711	0.84	12.88	0.6409	0.9736	0.29	11.62	0.5437	0.9874
DPO-Evader	0.06	9.82	0.4409	0.9206	0.31	6.80	0.3529	0.9271	0.14	8.37	0.5404	0.9329
HMGC	0.03	48.02	0.5181	0.9155	0.07	58.85	0.5705	0.9170	0.03	42.69	0.5283	0.9257
GradEscape	0.37	21.51	0.6151	0.9584	0.52	14.69	0.7015	0.9666	0.38	20.40	0.6162	0.9499
CoPA	0.20	27.52	0.6388	0.8727	0.19	29.95	0.6346	0.8740	0.16	23.75	0.6446	0.8698
Ours	0.87	20.67	0.6396	0.9048	0.98	17.54	0.7294	0.8993	1.00	20.08	0.6431	0.8194

Table 1: Evasion performance comparison against the fine-tuned RoBERTa detector (detection threshold $\tau = 0.5$).




Figure 4: Ablation study of MASH components, averaged across three detectors (RoBERTa, Binoculars, SCRN) and six datasets.

methods (e.g., DIPPER) demand substantial GPU memory. In contrast, MASH achieves a lightweight memory footprint (~ 3 GB) and rapid inference speed (1.7s). For a comprehensive comparison including training and Stage 4 overhead, Figure 12 in Appendix C shows that MASH amortizes its overhead effectively: as attack samples scale to 10^4 , the average query cost per sample becomes negligible.

4.3 Ablation Study

Impact of MASH Components. As shown in Figure 4, DPO significantly improves ASR, while the refinement module restores the text quality de-

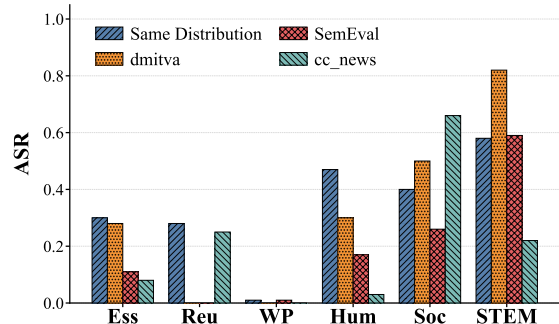


Figure 5: Impact of data source selection.




Figure 6: Impact of training data size.

graded by DPO. Detailed ablation results are provided in Table 9 (Appendix F).

Impact of Data Source and Size. We constructed the training data from four sources (see Figure 5). To establish a fair comparison, we downsampled open-source datasets to match the limited ‘‘Same Distribution’’ (i.e., matching the detector’s train-

Method	MGTBench-Essay				MGTBench-Reuters				MGTBench-WP			
	ASR ↑	PPL ↓	GRUEN ↑	BS ↑	ASR ↑	PPL ↓	GRUEN ↑	BS ↑	ASR ↑	PPL ↓	GRUEN ↑	BS ↑
DeepWordBug	0.00	11.46	0.5472	0.9767	0.00	11.42	0.6739	0.9796	0.00	12.40	0.6179	0.9823
TextBugger	0.00	12.61	0.5362	0.9644	0.00	13.65	0.6360	0.9578	0.00	14.43	0.6109	0.9706
TextFooler	0.00	12.86	0.5743	0.9781	0.00	13.64	0.6710	0.9746	0.00	15.09	0.6027	0.9760
DIPPER	0.23	9.97	0.6372	0.9041	0.16	11.08	0.7442	0.9006	0.33	12.63	0.6679	0.9086
Toblend	0.04	16.44	0.3931	0.7583	0.21	13.34	0.5293	0.7302	0.22	13.22	0.5303	0.7372
PromptAttack	0.53	10.00	0.7567	0.9169	0.55	9.98	0.8083	0.9327	0.35	10.30	0.6881	0.9264
Charmer	0.50	11.93	0.5505	0.9846	0.63	12.29	0.6677	0.9820	0.52	14.95	0.6233	0.9838
DPO-Evader	0.03	4.20	0.1305	0.9248	0.03	3.81	0.1641	0.9233	0.02	4.59	0.1591	0.9198
HMGC	0.32	13.67	0.5545	0.9781	0.41	13.00	0.6428	0.9779	0.31	14.26	0.6208	0.9773
GradEscape	0.64	14.45	0.6446	0.9558	0.62	13.27	0.7200	0.9594	0.37	12.29	0.6745	0.9683
CoPA	0.11	8.95	0.2516	0.9165	0.05	9.24	0.4468	0.9164	0.16	10.08	0.3921	0.9156
Ours	0.94	11.39	0.6755	0.9077	0.95	12.91	0.7122	0.9037	0.85	12.05	0.6883	0.9286

Method	MGT-Academic-Humanity				MGT-Academic-Social Science				MGT-Academic-STEM			
	ASR ↑	PPL ↓	GRUEN ↑	BS ↑	ASR ↑	PPL ↓	GRUEN ↑	BS ↑	ASR ↑	PPL ↓	GRUEN ↑	BS ↑
DeepWordBug	0.00	10.73	0.6387	0.9822	0.00	10.85	0.6050	0.9828	0.00	8.43	0.5746	0.9871
TextBugger	0.00	11.81	0.6173	0.9675	0.00	13.39	0.5854	0.9628	0.00	10.42	0.5485	0.9680
TextFooler	0.00	11.37	0.6369	0.9813	0.00	13.38	0.6143	0.9794	0.00	9.92	0.5776	0.9840
DIPPER	0.53	11.93	0.6378	0.8917	0.38	11.63	0.6407	0.8891	0.49	10.86	0.6429	0.9048
PromptAttack	0.83	11.13	0.7382	0.9270	0.86	10.67	0.7320	0.9276	0.81	9.33	0.6904	0.9318
Charmer	0.71	12.91	0.6304	0.9831	0.77	12.37	0.5960	0.9832	0.71	10.16	0.5592	0.9841
DPO-Evader	0.08	5.79	0.2598	0.9502	0.06	5.17	0.1590	0.9440	0.04	5.02	0.1795	0.9525
HMGC	0.69	15.05	0.6024	0.9732	0.68	19.14	0.5847	0.9715	0.73	15.25	0.5384	0.9708
GradEscape	0.87	17.58	0.7000	0.9416	0.88	16.86	0.6649	0.9422	0.95	20.92	0.6546	0.9286
CoPA	0.25	9.25	0.3889	0.9273	0.17	9.52	0.3438	0.9258	0.23	7.85	0.3377	0.9328
Ours	0.95	13.47	0.7025	0.9074	0.95	13.31	0.7014	0.9073	0.99	14.02	0.6434	0.8864

Table 2: Evasion performance comparison against the Binoculars detector.

Detector	Dataset	ASR	BS	No Attack		With Attack	
				PPL	GRUEN	PPL	GRUEN
Scribbr	Essay	0.95	0.8800	7.33	0.5903	16.89	0.6599
	Reuters	0.59	0.8740	8.87	0.6964	8.54	0.7318
	WP	0.90	0.8844	8.18	0.7133	17.18	0.6302
	Humanity	0.94	0.9288	10.82	0.5899	19.61	0.6712
	Social	0.98	0.8730	9.50	0.6714	18.53	0.7333
	STEM	1.00	0.7840	8.61	0.5690	20.29	0.6442
Writer	Essay	0.90	0.8722	7.41	0.5727	15.87	0.6806
	Reuters	0.66	0.8624	7.91	0.6440	9.16	0.7440
	WP	0.91	0.8947	8.79	0.6524	15.78	0.6944
	Humanity	0.86	0.9067	8.39	0.6511	20.32	0.6246
	Social	0.96	0.9062	7.60	0.5983	16.42	0.7092
	STEM	0.99	0.7832	7.91	0.6102	18.29	0.6741

Table 3: Commercial detector evasion.

ing domain) baseline counts. Experimental results reveal that the effectiveness of generic datasets varies substantially across domains. Notably, for the RoBERTa detector trained on Reuters, samples from *dmitva*¹ and SemEval (Wang et al., 2024b) were predominantly classified as machine-generated, rendering them unsuitable for initializing style transfer in this domain. In contrast, *cc_news*² proved effective. Figure 6 further confirms that performance scales with data size.

Generalizability of Refinement Module. To assess whether Stage 4 generalizes beyond MASH,

¹https://huggingface.co/datasets/dmitva/human_ai_generated_text

²https://huggingface.co/datasets/vblagoje/cc_news

Method	GPU Memory	Inference Time	Query Cost
HMGC	41046	86	0
Charmer	30094	189.9	18827
DPO_Evader	22699	0.2	0
DIPPER	22241	58.6	0
CoPA	15425	6.7	0
ToBlend	9223	31.8	0
TextBugger	3753	78.3	5275.7
TextFooler	3753	40.6	2913.8
GradEscape	3133	1.7	0
DeepWordBug	1697	29.1	1018.4
MASH	3139	1.7	0

Table 4: Comparison of resource usage. Units: GPU Memory (MiB), Inference Time (seconds/sample), and Query Cost (queries/sample).

we apply it to outputs generated by other baseline attack methods. As shown in Figure 7, the refinement module consistently improves text quality across most baselines. Detailed results are provided in Table 10 (Appendix F).

Alternative Refinement Models. To reduce the cost of Stage 4, we evaluate open-source alternatives. As shown in Table 8 (Appendix F), Qwen2.5-7B-Instruct improves GRUEN from 0.468 to 0.631 on average, approaching GPT-5’s 0.667. However, GPT-5 maintains a clear advantage in BERTScore (0.887 vs. 0.857). These open-source alternatives remain viable under budget constraints.




Figure 7: Generalizability of inference-time refinement (Stage 4) to baseline attack methods, averaged across six datasets. Gray and red bars denote original and refined outputs, respectively; error bars indicate standard deviation.




Figure 8: Cross-domain evasion transferability.

4.4 Universality and Transferability Analysis

Transferability across Domains. We evaluate cross-domain transferability by testing MASH against RoBERTa detectors fine-tuned on source domains and evaluated on different target domains. As illustrated in Figure 8, the results demonstrate robust generalization across most datasets. However, cross-domain scenarios involving Reuters exhibit an ASR approaching zero. We provide a detailed analysis of this phenomenon in Appendix G.

Transferability across Detectors. To assess whether MASH captures universal human-like characteristics or merely overfits specific detector boundaries, we evaluate cross-detector transferability. We additionally introduce SCRN (Huang et al., 2024b), another RoBERTa-based fine-tuned detector. As shown in Figure 9, attacks optimized against supervised detectors (RoBERTa, SCRN) transfer effectively to zero-shot methods (Binoculars), whereas the reverse fails. This asymmetry suggests that supervised detector optimization implicitly corrects statistical artifacts exploited by zero-shot methods.




Figure 9: Cross-detector evasion transferability.

4.5 Defense via Adversarial Training

To mitigate risks, we explore adversarial training by fine-tuning detectors on MASH-polished samples. This defense reduces ASR from 92% to near-zero, but degrades accuracy on clean samples by 18.7% on average. Details are provided in Appendix D.

5 Conclusion

We introduce MASH, a black-box framework that reformulates detector evasion as a style transfer task. Experiments show MASH outperforms state-

of-the-art baselines against both open-source and commercial detectors, exposing the fragility of current defenses. MASH serves as both a red-teaming tool and a foundation for more robust detection.

Limitations

Limited scope of the experiments. While we have conducted a comprehensive evaluation across five representative detectors, six diverse domains, and eleven state-of-the-art baseline methods, the landscape of AI-generated text detection is rapidly evolving. We acknowledge that our benchmarking represents a snapshot of the current security landscape. Future work could extend the MASH framework to test against emerging detection paradigms and novel adversarial attack strategies, further validating its generalization capabilities in dynamic adversarial environments.

Prerequisites for Distribution Alignment. The effectiveness of our Style-SFT module relies on the availability of a seed set of human-written texts that theoretically fall within the target detector’s “human” decision boundary. As indicated in our ablation study (see Figure 5), significant distribution shifts in generic open-source corpora can hinder initialization. However, this dependency is not a methodological flaw but a logical prerequisite: evasion is only mathematically meaningful against detectors that maintain a reasonable False Positive Rate (FPR). If a detector fails to correctly classify ground-truth human text (i.e., exhibits high FPR), the adversarial objective of “mimicking human style” becomes redundant.

Linguistic Generalization. Our current evaluations are concentrated on English-language benchmarks, including MGTBench and MGT-Academic. The applicability of our style-injection and DPO alignment mechanisms to languages with complex morphological structures or low-resource settings remains unexplored. Future work could investigate the cross-lingual transferability of MASH to assess its robustness in diverse linguistic contexts.

Ethics Statement

This research investigates the vulnerabilities of current AI-generated text detectors under black-box settings. The primary motivation for proposing MASH is to serve as a red-teaming framework to expose the fragility of existing detection paradigms. We acknowledge that the techniques described herein could potentially be misused for academic

dishonesty or disinformation dissemination. However, we believe that security through obscurity is unsustainable; publicly disclosing these vulnerabilities is essential for the community to develop more robust detection systems.

Furthermore, all experiments in this work utilize open-source datasets and models, strictly adhering to their usage policies and data privacy regulations. We urge developers of detection systems to move beyond simple statistical features and incorporate adversarial training to defend against evasion via style transfer.

Acknowledgment

This work is supported in part by New Generation Artificial Intelligence-National Science and Technology Major Project (2025ZD0123504).

References

- Elias Abad Rocamora, Yongtao Wu, Fanghui Liu, Grigorios G Chrysos, and Volkan Cevher. 2024. Revisiting character-level adversarial attacks for language models. In *International Conference on Machine Learning (ICML)*.
- Jinheon Baek, Sujay Kumar Jauhar, Silviu Cucerzan, and Sung Ju Hwang. 2025. Researchagent: Iterative research idea generation over scientific literature with large language models. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6709–6738.
- Guangsheng Bao, Yanbin Zhao, Zhiyang Teng, Linyi Yang, and Yue Zhang. 2023. Fast-detectgpt: Efficient zero-shot detection of machine-generated text via conditional probability curvature. *arXiv preprint arXiv:2310.05130*.
- Ralph Allan Bradley and Milton E Terry. 1952. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345.
- Zihao Cheng, Li Zhou, Feng Jiang, Benyou Wang, and Haizhou Li. 2025. Beyond binary: Towards fine-grained llm-generated text detection via role recognition and involvement measurement. In *Proceedings of the ACM on Web Conference 2025*, pages 2677–2688.
- Hao Fang, Jiawei Kong, Tianqu Zhuang, Yixiang Qiu, Kuofeng Gao, Bin Chen, Shu-Tao Xia, Yaowei Wang, and Min Zhang. 2025. [Your language model can secretly write like humans: Contrastive paraphrase attacks on LLM-generated text detectors](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 8596–8613,

- Suzhou, China. Association for Computational Linguistics.
- Ji Gao, Jack Lanchantin, Mary Lou Soffa, and Yanjun Qi. 2018. Black-box generation of adversarial text sequences to evade deep learning classifiers. In *2018 IEEE Security and Privacy Workshops (SPW)*, pages 50–56. IEEE.
- Sebastian Gehrmann, Hendrik Strobelt, and Alexander Rush. 2019. *GLTR: Statistical detection and visualization of generated text*. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 111–116, Florence, Italy. Association for Computational Linguistics.
- Xun Guo, Yongxin He, Shan Zhang, Ting Zhang, Wanguan Feng, Haibin Huang, and Chongyang Ma. 2024. Detective: Detecting ai-generated text via multi-level contrastive learning. *Advances in Neural Information Processing Systems*, 37:88320–88347.
- Abhimanyu Hans, Avi Schwarzschild, Valeriia Cherepanova, Hamid Kazemi, Aniruddha Saha, Micah Goldblum, Jonas Geiping, and Tom Goldstein. 2024. Spotting llms with binoculars: Zero-shot detection of machine-generated text. *arXiv preprint arXiv:2401.12070*.
- Xinlei He, Xinyue Shen, Zeyuan Chen, Michael Backes, and Yang Zhang. 2024. Mgtbench: Benchmarking machine-generated text detection. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, pages 2251–2265.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P Xing. 2017. Toward controlled generation of text. In *International conference on machine learning*, pages 1587–1596. PMLR.
- Fan Huang, Haewoon Kwak, and Jisun An. 2024a. To-blend: Token-level blending with an ensemble of llms to attack ai-generated text detection. *arXiv preprint arXiv:2402.11167*.
- Guanhua Huang, Yuchen Zhang, Zhe Li, Yongjian You, Mingze Wang, and Zhouwang Yang. 2024b. *Are AI-generated text detectors robust to adversarial perturbations?* In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6005–6024, Bangkok, Thailand. Association for Computational Linguistics.
- Lei Jiang, Desheng Wu, and Xiaolong Zheng. 2025. Sendetex: Sentence-level ai-generated text detection for human-ai hybrid content via style and context fusion. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 5287–5302.
- JianWang, Shangqing Liu, Xiaofei Xie, and Yi Li. 2024. An empirical study to evaluate aigc detectors on code content. In *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering*, pages 844–856.
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2020. Is bert really robust? a strong baseline for natural language attack on text classification and entailment. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 8018–8025.
- Kalpesh Krishna, Yixiao Song, Marzena Karpinska, John Wieting, and Mohit Iyyer. 2023. Paraphrasing evades detectors of ai-generated text, but retrieval is an effective defense. *Advances in Neural Information Processing Systems*, 36:27469–27500.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 7871–7880.
- Jia Li, Chongyang Tao, Jia Li, Ge Li, Zhi Jin, Huangzhao Zhang, Zheng Fang, and Fang Liu. 2025. Large language model-aware in-context learning for code generation. *ACM Transactions on Software Engineering and Methodology*, 34(7):1–33.
- Jinfeng Li, Shouling Ji, Tianyu Du, Bo Li, and Ting Wang. 2018. Textbugger: Generating adversarial text against real-world applications. *arXiv preprint arXiv:1812.05271*.
- Yafu Li, Qintong Li, Leyang Cui, Wei Bi, Zhilin Wang, Longyue Wang, Linyi Yang, Shuming Shi, and Yue Zhang. 2024. Mage: Machine-generated text detection in the wild. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 36–53.
- Dongqi Liu and Vera Demberg. 2023. Chatgpt vs human-authored text: Insights into controllable text summarization and sentence style transfer. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 4: Student Research Workshop)*, pages 1–18.
- Yule Liu, Zhiyuan Zhong, Yifan Liao, Zhen Sun, Jingyi Zheng, Jiaheng Wei, Qingyuan Gong, Fenghua Tong, Yang Chen, Yang Zhang, and Xinlei He. 2025. On the generalization and adaptation ability of machine-generated text detectors in academic writing. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pages 5674–5685.
- Ning Lu, Shengcai Liu, Rui He, Yew-Soon Ong, Qi Wang, and Ke Tang. 2024. *Large language models can be guided to evade AI-generated text detection*. *Transactions on Machine Learning Research*.
- Dominik Macko, Jakub Kopal, Robert Moro, and Ivan Srba. 2025. Multisocial: Multilingual benchmark of machine-generated text detection of social-media texts. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 727–752.

- Chengzhi Mao, Carl Vondrick, Hao Wang, and Junfeng Yang. 2024. [Raidar: generative AI detection via rewriting](#). In *The Twelfth International Conference on Learning Representations*.
- Wenlong Meng, Shuguo Fan, Chengkun Wei, Min Chen, Yuwei Li, Yuanchao Zhang, Zhikun Zhang, and Wenzhi Chen. 2025. Gradscape: a gradient-based evader against ai-generated text detectors. SEC '25, USA. USENIX Association.
- Eric Mitchell, Yoonho Lee, Alexander Khazatsky, Christopher D Manning, and Chelsea Finn. 2023. Detectgpt: Zero-shot machine-generated text detection using probability curvature. In *International conference on machine learning*, pages 24950–24962. PMLR.
- Charlotte Nicks, Eric Mitchell, Rafael Rafailov, Archit Sharma, Christopher D Manning, Chelsea Finn, and Stefano Ermon. 2023. Language model detectors are easily optimized against. In *The twelfth international conference on learning representations*.
- Andrea Pedrotti, Michele Papucci, Cristiano Ciaccio, Alessio Miaschi, Giovanni Puccetti, Felice Dell’Orletta, and Andrea Esuli. 2025. [Stress-testing machine generated text detection: Shifting language models writing style to fool detectors](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 3010–3031, Vienna, Austria. Association for Computational Linguistics.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36:53728–53741.
- Emily Reif, Daphne Ippolito, Ann Yuan, Andy Coenen, Chris Callison-Burch, and Jason Wei. 2022. A recipe for arbitrary text style transfer with large language models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 837–848.
- Jenna Russell, Marzena Karpinska, and Mohit Iyyer. 2025. [People who frequently use ChatGPT for writing tasks are accurate and robust detectors of AI-generated text](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5342–5373, Vienna, Austria. Association for Computational Linguistics.
- Vinu Sankar Sadasivan, Aounon Kumar, Sriram Balasubramanian, Wenxiao Wang, and Soheil Feizi. 2023. Can ai-generated text be reliably detected? *arXiv preprint arXiv:2303.11156*.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. Style transfer from non-parallel text by cross-alignment. *Advances in neural information processing systems*, 30.
- Tianchun Wang, Yuanzhou Chen, Zichuan Liu, Zhanwen Chen, Haifeng Chen, Xiang Zhang, and Wei Cheng. 2025. [Humanizing the machine: Proxy attacks to mislead LLM detectors](#). In *The Thirteenth International Conference on Learning Representations*.
- Yichen Wang, Shangbin Feng, Abe Hou, Xiao Pu, Chao Shen, Xiaoming Liu, Yulia Tsvetkov, and Tianxing He. 2024a. Stumbling blocks: Stress testing the robustness of machine-generated text detectors under attacks. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2894–2925.
- Yuxia Wang, Jonibek Mansurov, Petar Ivanov, Jinyan Su, Artem Shelmanov, Akim Tsvigun, Osama Mohammed Afzal, Tarek Mahmoud, Giovanni Puccetti, Thomas Arnold, Chenxi Whitehouse, Alham Fikri Aji, Nizar Habash, Iryna Gurevych, and Preslav Nakov. 2024b. Semeval-2024 task 8: Multidomain, multimodel and multilingual machine-generated text detection. *arXiv preprint arXiv:2404.14183*.
- Xilie Xu, Keyi Kong, Ning Liu, Lizhen Cui, Di Wang, Jingfeng Zhang, and Mohan Kankanhalli. 2024. [An LLM can fool itself: A prompt-based adversarial attack](#). In *The Twelfth International Conference on Learning Representations*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.
- Jingyi Zheng, Junfeng Wang, Zhen Sun, Wenhan Dong, Yule Liu, and Xinlei He. 2025. Th-bench: Evaluating evading attacks via humanizing ai text on machine-generated text detectors. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pages 5948–5959.
- Ying Zhou, Ben He, and Le Sun. 2024. Humanizing machine-generated content: Evading ai-text detection through adversarial attack. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation*.
- Wanzheng Zhu and Suma Bhat. 2020. Gruen for evaluating linguistic quality of generated text. *arXiv preprint arXiv:2010.02498*.

A Theoretical Analysis

This section serves as a theoretical supplement to the methodology presented in Section 3.2. We provide formal justifications for the MASH framework, specifically elucidating the mechanism of DPO-based evasion, the construction of preference pairs, and the necessity of the multi-stage pipeline.

A.1 Distribution Alignment via DPO

To theoretically ground our approach, we bridge the gap between the adversarial evasion goal and the DPO optimization objective. In our framework, the preference labels are derived from a black-box detector $D(\mathbf{y}) \in [0, 1]$, which estimates the probability of \mathbf{y} being AI-generated. We posit that the underlying human preference follows a Bradley-Terry model driven by an implicit reward function $r(\mathbf{x}, \mathbf{y})$. Following Wang et al. (2025), we formulate this reward as a scaled and shifted version of the detector’s confidence score. For theoretical convenience and without loss of generality, we align the reward definition in Eq. 6 with the Bradley-Terry preference model by adopting the following equivalent form:

$$r(\mathbf{x}, \mathbf{y}) = C \cdot (1 - D(\mathbf{y})), \quad (9)$$

where C is a sufficiently large constant. Under this formulation, maximizing the implicit reward is mathematically equivalent to minimizing the detector’s AI probability score. Crucially, the large constant C renders the Bradley-Terry probability nearly deterministic, thereby theoretically aligning the DPO optimization objective with our adversarial evasion goal.

Theorem 1 (Optimality of Evasion). *Let π_{ref} be the reference policy and π^* be the optimal policy minimizing the DPO loss. The generated distribution π^* aligns with the human-like regions defined by the detector D (i.e., regions where $D(\mathbf{y}) \rightarrow 0$).*

Proof. As derived by Rafailov et al. (2023), the optimal policy π^* for the KL-constrained reward maximization problem is uniquely determined by the reward function and the reference policy, taking the following closed-form solution:

$$\pi^*(\mathbf{y}|\mathbf{x}) = \frac{1}{Z(\mathbf{x})} \pi_{\text{ref}}(\mathbf{y}|\mathbf{x}) \exp\left(\frac{1}{\beta} r(\mathbf{x}, \mathbf{y})\right), \quad (10)$$

where $Z(\mathbf{x})$ is the partition function. Substituting our evasion-oriented implicit reward $r \propto (1 - D(\mathbf{y}))$ into Eq. 10:

$$\pi^*(\mathbf{y}|\mathbf{x}) \propto \pi_{\text{ref}}(\mathbf{y}|\mathbf{x}) \exp\left(\frac{C}{\beta} (1 - D(\mathbf{y}))\right). \quad (11)$$

As the coefficient $\frac{C}{\beta}$ is large, the exponential term acts as a sharp filter. It amplifies the probability mass of samples where $1 - D(\mathbf{y}) \approx 1$ (i.e., $D(\mathbf{y}) \approx 0$, corresponding to human-written text)

	Essay	Reuters	WP	Humanity	Social
Stage 1	0.83	0.09	0.68	0.68	0.92
Stage 2 (AP) [†]	0.87	0.61	0.84	0.66	0.87
Stage 2 (HN) [‡]	0.95	0.73	0.9	0.87	0.98

Note: AP = Ambiguous Pairs, HN = Hard Negatives.

Table 5: Comparison of ASR across different domains.

and suppresses regions where $D(\mathbf{y}) \approx 1$. Consequently, minimizing the DPO loss forces the generator’s distribution to converge towards the human distribution characterized by the detector. \square

A.2 Analysis of Hard Negatives

We theoretically justify why constructing hard negative pairs is superior to random or easy negatives.

Proposition 1 (Optimization Consistency & Gradient Efficiency). *Constructing preference pairs with a maximal detector score gap $\Delta D = D(y_l) - D(y_w) \approx 1$ is necessary to ensure consistency between empirical labels and the Bradley-Terry model, and to maximize the optimization gradient.*

Proof. We define the evasion-oriented implicit reward as $r(y) = -C \cdot D(y)$ (where $C > 0$), meaning higher detection scores yield lower rewards. Substituting this into the Bradley-Terry model:

$$\begin{aligned} P_{\text{BT}}(y_w \succ y_l) &= \sigma(r(y_w) - r(y_l)) \\ &= \sigma(-CD_w - (-CD_l)) \\ &= \sigma\left(C \cdot \underbrace{(D_l - D_w)}_{\Delta D}\right), \end{aligned} \quad (12)$$

where σ denotes the logistic sigmoid function. This derivation confirms that maximizing the probability requires maximizing the gap ΔD . Table 5 confirms that employing hard negative samples yields the most significant improvement in ASR.

- **Ambiguous Pairs ($\Delta D \approx 0$):** If y_l evades detection ($D_l \approx 0$), then $\Delta D \approx 0$ and $P_{\text{BT}} \approx \sigma(0) = 0.5$. This contradicts the deterministic hard label ($y_w \succ y_l$), introducing label noise.
- **Hard Negatives ($\Delta D \approx 1$):** With $D_l \approx 1$ and $D_w \approx 0$, we have $P_{\text{BT}} \approx \sigma(C) \rightarrow 1$. This aligns with the empirical label.

The gradient of the DPO loss \mathcal{L}_{DPO} in Eq. 8 with respect to model parameters θ is explicitly scaled

by a weighting term $w(\theta)$:

$$\nabla_{\theta} \mathcal{L}_{\text{DPO}} = -\mathbb{E} \left[\underbrace{\sigma(\hat{r}_l - \hat{r}_w)}_{w(\theta)} \nabla_{\theta} \log \frac{\pi_{\theta}(y_w|x)}{\pi_{\theta}(y_l|x)} \right], \quad (13)$$

where \hat{r}_l and \hat{r}_w denote the implicit rewards $\hat{r}_{\theta}(y_l)$ and $\hat{r}_{\theta}(y_w)$, respectively. For hard negatives, the model initially assigns high probability to the machine artifact y_l (i.e., it confidently errs), leading to an inverted reward estimate where $\hat{r}_l > \hat{r}_w$. Consequently, the weighting term $w(\theta) > 0.5$, generates a strong error-correction signal. Conversely, easy negatives result in $\hat{r}_l \ll \hat{r}_w$, causing $w(\theta) \rightarrow 0$ and leading to vanishing gradients. \square

A.3 The Necessity of Style-SFT

Finally, we prove the necessity of Style-SFT (Stage 2) for effective DPO alignment.

Theorem 2 (Support Constraint). *Let $\mathcal{Y}_H = \{\mathbf{y} \mid D(\mathbf{y}) < \tau\}$ denote the target human-like region. If the reference policy π_{ref} assigns negligible mass to \mathcal{Y}_H , the DPO-optimized policy π^* cannot converge to \mathcal{Y}_H regardless of the reward magnitude.*

Proof. Recalling the closed-form optimal policy derived in Eq. 10. This multiplicative relationship implies that π^* is strictly bounded by the support of π_{ref} . We analyze two initialization scenarios:

- **Weak Initialization (Zero Support):** Without style injection, standard LLMs often exhibit a domain gap, such that $\forall \mathbf{y} \in \mathcal{Y}_H, \pi_{\text{ref}}(\mathbf{y}|\mathbf{x}) \approx 0$. Consequently, even if the implicit reward r is maximized for human text, the posterior probability $\pi^*(\mathbf{y}|\mathbf{x}) \approx 0 \cdot \exp(\dots) = 0$. The model faces a “cold start” problem where high-reward regions are theoretically unreachable.
- **Style-SFT Initialization (Warm Start):** Stage 2 explicitly shifts probability mass towards the target style, ensuring $\pi_{\text{SFT}}(\mathbf{y}|\mathbf{x}) > \epsilon$, where ϵ is a minimal existence threshold. Crucially, a larger initial probability π_{SFT} minimizes the distribution shift required to reach π^* , thereby reducing the optimization effort needed to satisfy the exponential scaling.

Thus, Stage 2 is a mathematical prerequisite to ensure the target region lies within the feasible search space of DPO. \square

B Experimental Details

Datasets. We conduct experiments on two diverse benchmarks: (1) MGTBench (He et al., 2024), which comprises diverse writing styles: Essays (student argumentations), WP (creative fiction from Reddit WritingPrompts), and Reuters (financial and global news). Its AI samples are generated by seven LLMs (e.g., GPT-3.5, Llama, and Claude) using task-specific prompts to mimic human writing. (2) MGT-Academic (Liu et al., 2025), covering formal scholarly discourse across: STEM (scientific and technical papers from ArXiv), Social Science (encyclopedia entries from Wikipedia), and Humanity (literary books from Project Gutenberg). Advanced LLMs (e.g., GPT-4, Gemini) generate these counterparts via academic paraphrasing or completion while preserving formal discourse.

To ensure rigorous evaluation, each subset is strictly partitioned into four splits—detector training, validation, shadow training, and testing—with a ratio of 3:1:1:1. The first two are used for fine-tuning the target RoBERTa detector, while the shadow set serves to train surrogate models for white-box baselines. Following Fang et al. (2025), we randomly sample 100 machine-generated instances per domain for the final testing.

Target Detectors. We evaluate evasion performance against five black-box detectors. (1) Open-source: RoBERTa is fine-tuned on the training split of each subset, serving as a robust supervised baseline as identified by Zheng et al. (2025). Binoculars (Hans et al., 2024) is a zero-shot detector requiring no additional training. SCRIN (Huang et al., 2024b) is used with its official pre-trained weights. (2) Commercial: Following JianWang et al. (2024), we include Writer³ and Scribbr⁴. All detectors are evaluated under black-box settings, with access limited to confidence scores or binary labels.

Baselines. We benchmark MASH against 11 state-of-the-art evasion methods, categorized into perturbation-, prompt-, and paraphrase-based attacks. Perturbation-based methods include DeepWordBug, TextBugger, TextFooler, Charmer, and HMGC. PromptAttack is the prompt-based baseline, which guides LLMs to generate evasion-prone text through optimized instructions. Paraphrase-based methods comprise DIPPER, Toblend, CoPA, GradEscape, and DPO-Evader. To maintain a strict black-box setting, HMGC and GradEscape are im-

³<https://writer.com/ai-content-detector/>

⁴<https://www.scribbr.com/ai-detector/>

plemented by training auxiliary RoBERTa detectors on shadow training data.

For DIPPER, we configure both lexical and order diversity to 40. In the Toblend evaluation, we employ GPT2-XL and OPT-2.7B as the source generators. Due to the absence of original generation prompts in the MGT-Academic benchmark, Toblend’s assessment is restricted to the MGTBench dataset. CoPA is implemented using Qwen-7B as the backbone paraphraser. For DPO-Evader, we adopt Qwen-0.5B as the paraphraser to accommodate GPU memory constraints during training. Empirically, we observed that the 0.5B model achieves higher attack success rates than larger variants (e.g., 3B+), suggesting that smaller models may introduce more textual noise conducive to distribution alignment. All baseline implementations and hyperparameter configurations follow their respective original protocols unless otherwise specified.

Metrics. Attack Success Rate (ASR) is the proportion of polished texts $\{x_{adv}^{(i)}\}_{i=1}^N$ classified as “human-written” by detector D :

$$ASR = \frac{1}{N} \sum_{i=1}^N \mathbb{1}(D(x_{adv}^{(i)}) < \tau), \quad (14)$$

where $\mathbb{1}(\cdot)$ is the indicator function and τ is the decision threshold.

Following Zhou et al. (2024), we assess fluency via Perplexity (PPL), calculated using Pythia-2.8b. A lower PPL indicates superior text fluency. For a sequence x of length T , PPL is computed as the exponential of the average negative log-likelihood:

$$PPL(x) = \exp\left(-\frac{1}{T} \sum_{t=1}^T \log P(x_t | x_{<t})\right). \quad (15)$$

To evaluate semantic consistency between x_{ai} and x_{adv} , we employ BERTScore (Zhang et al., 2019). BERTScore computes token-level similarity using contextual embeddings and aggregates them into precision, recall, and F1 metrics. Following standard practice, we report the F1-score to provide a robust assessment of semantic integrity.

Distinct from PPL, we employ GRUEN (Zhu and Bhat, 2020) for a holistic, reference-free linguistic quality assessment. Higher GRUEN values indicate better overall linguistic quality.

Implementation and Training Details. All experiments were conducted on a single NVIDIA RTX 3090 GPU, with a complete training cycle per domain lasting approximately 7.5 hours.

During the Style-SFT stage, We jointly optimize \mathcal{L}_{recon} and \mathcal{L}_{trans} with $\lambda = 0.5$. We employ the AdamW optimizer with a learning rate of 2×10^{-5} and a batch size of 8. The model is trained for 50 epochs with early stopping.

During the DPO phase, the learning rate is adjusted to 5×10^{-6} for 5 epochs. To accommodate GPU memory constraints, we use a per-device batch size of 2 combined with 8 gradient accumulation steps. For inference, we utilize beam search with a beam size of 4.

For the Reuters domain in MGTBench, we utilize the vblagoje/cc_news dataset, while for all other domains, we employ the dmitva/human_ai_generated_text dataset. We paraphrase human sources using Qwen2.5-3B-Instruct, Llama-3-8B-Instruct, and GPT-2. The vLLM framework is adopted to accelerate the inference process. This pipeline yielded approximately 6,000 validated pairs, achieving a label-flip success rate (transitioning from human-written to AI-generated) of approximately 95%.

Inverse Data Construction (Stage 1)

System Prompt: You are a professional English editor focusing on clarity and conciseness.

User Prompt: Polish the following English text for grammar, clarity, and flow. Output only the polished text — no introductory phrases, no explanations, no headings, and no markdown or asterisk * lists. Preserve any in-text citations exactly as they appear: \n\n {human_text}

In Stage 4, sentence-level polish candidates are generated by querying the GPT-5 API. This refinement module ensures grammatical fluency while strictly preserving the evasion status of the text.

C Additional Experimental Results

Visualization of Decision Boundaries. The superior performance of current AI-generated text detectors is fundamentally rooted in the distinct distribution discrepancy between machine-generated and human-written text. As shown in Figure 10, based on 1,200 samples from the Essay dataset, distinct separation is observable in both the deep semantic feature space extracted by supervised detectors and the statistical feature space utilized by zero-shot detectors. Consequently, we posit that effective evasion should not be merely adversarial noise injection against a detector, but rather a style transfer task.

Token-Level Detection Contribution. We hypothesize that the effectiveness of perturbation-

Sentence-level Refinement (Stage 4)

System Prompt: You are a strict Grammar Correction and Semantic Alignment Specialist. Your goal is to refine the text to be grammatically perfect and fluent, while ensuring the meaning aligns rigorously with the provided ground truth.

YOUR TASKS

1. Grammar & Fluency: Correct grammatical errors, or unnatural expressions in the input sentences.
2. Semantic Alignment (Crucial): Compare each input sentence against the 'Reference Context'. If an input sentence is vague, misleading, or deviates from the Reference Context, rewrite it to strictly match the Reference Context's meaning.
3. Structure Flexibility: Replace or delete sentences to improve clarity and semantic consistency. Ensure the final output maintains a logical sequence.

OUTPUT FORMAT

Output ONLY a JSON string array of arrays (List[List[str]]):

1. Output MUST be a nested array of {input_len} elements.
2. Do NOT explain your changes. Output JSON only.

EXAMPLE

Reference: "Photosynthesis allows plants to convert sunlight into energy. This process creates oxygen, which is vital for life on Earth."

Input List: ["Plants use sun to make power and this creates air for everyone.", "This is a very cool thing.", "It is important for the world."]

Output: [["Photosynthesis allows plants to convert sunlight into energy.", "This process also generates oxygen."], [], ["This function is vital for sustaining life on Earth."]]

User Prompt: ### Reference Context

{reference_text}

Input List

{sentences_list}

based attacks correlates with token contribution density. For validation, we analyzed 100 randomly sampled instances from MultiSocial (Macko et al., 2025) and MGTBench. As illustrated in Figure 11, when high-contribution tokens are concentrated, perturbation-based attacks achieve a high hit rate, easily flipping classifications through local edits. However, when the decision influence is sparsely distributed across a vast "neutral zone," these attacks predominantly target non-critical regions, failing to induce the substantive shift in global feature distribution necessary to evade robust detectors.

Commercial Detector Baselines. To establish a rigorous benchmark, we assessed the detection capabilities of Writer and Scribbr across all six domains. As shown in Table 6, we report four standard classification metrics: Precision (the proportion of detected AI texts that are truly AI-generated), Recall (the proportion of AI-generated texts correctly identified), F1 Score (the harmonic mean of Precision and Recall), and Accuracy (the proportion of correctly classified samples). Both systems exhibit high precision, with Writer demonstrating superior recall and achieving 0.99 accuracy on the Essay dataset. These results confirm their

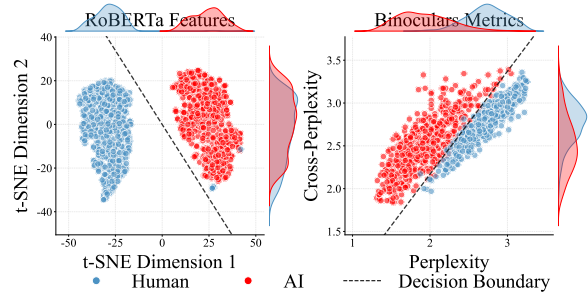


Figure 10: Visualization of the feature distribution gap between AI-generated and human-written text. The left plot shows the t-SNE projection of semantic features from the RoBERTa detector. The right plot illustrates the statistical features from the Binoculars detector. Marginal KDEs highlight dense regions.

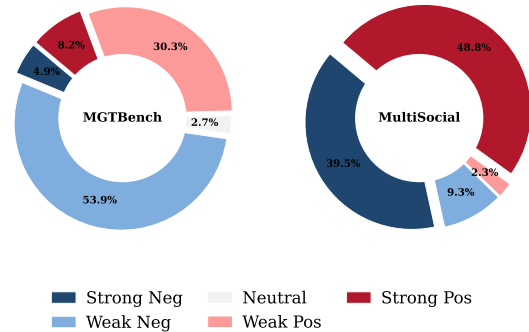


Figure 11: Attribution distribution of token contributions on MultiSocial and MGTBench, calculated via gradient-weighted embedding attribution. The left shows sparsely distributed decision influence; the right shows concentrated high-contribution tokens.

validity as robust baselines for evaluating evasion attacks. The total query cost for all commercial detector experiments was under \$20.

Amortized Query Cost Analysis. Figure 12 shows the average query cost per sample as a function of the total attack volume (N). Perturbation-based methods (e.g., Charmer, TextFooler) require iterative optimization for each input, resulting in a constant and high per-sample query cost regardless of scale. In contrast, MASH incurs a fixed upfront cost for training the paraphraser (Stages 2 and 3), but requires no additional queries during inference. As N increases, this training cost is distributed across more samples, leading to a sharp decrease in per-sample overhead. For large-scale attacks ($N > 10^4$), the average cost of MASH becomes negligible, achieving orders-of-magnitude better efficiency than query-based methods. This confirms that MASH is highly cost-effective for

Dataset	Precision	Recall	F1 Score	Accuracy
<i>Target Detector: Writer</i>				
Essay	1.0000	0.9800	0.9899	0.9900
Reuters	1.0000	0.8571	0.9231	0.9293
WP	0.9800	0.9800	0.9800	0.9800
Humanities	0.9091	0.6122	0.7317	0.7778
Social	0.9697	0.6400	0.7711	0.8100
STEM	1.0000	0.5800	0.7342	0.7900
<i>Target Detector: Scribbr</i>				
Essay	1.0000	0.5400	0.7013	0.7700
Reuters	1.0000	0.3600	0.5294	0.6800
WP	1.0000	0.5000	0.6667	0.7500
Humanities	1.0000	0.3400	0.5075	0.6700
Social	1.0000	0.4200	0.5915	0.7100
STEM	1.0000	0.4800	0.6486	0.7400

Table 6: Performance metrics of commercial detectors.




Figure 12: Amortized query cost per sample. MASH’s initial training overhead is rapidly diluted as attack volume grows, achieving significant efficiency gains over perturbation-based methods at scale.

high-volume evasion scenarios.

D Defensive Analysis

While MASH effectively evades existing detectors, we further investigate potential mitigation strategies against such style-based humanization attacks. We propose an adversarial training approach in which the target detector is fine-tuned on a mixture of original human texts and MASH-polished adversarial examples. We use RoBERTa as the base detector and construct a balanced training set of 1,000 human-written texts and 1,000 MASH-rewritten machine-generated texts.

The results in Table 7 reveal a clear robustness-accuracy trade-off. Adversarial training substantially improves the detector’s robustness against MASH: the ASR drops to nearly 0.00 across most domains (e.g., STEM, Reuters), demonstrating that the detector can learn to recognize humanized styles. However, this defense improves detection

Dataset	w/o Training		w Training	
	Accuracy	ASR	Accuracy	ASR
Essay	0.9974	0.95	0.3059	0.01
Reuters	0.9988	0.73	0.9988	0
WP	0.9912	0.90	0.8731	0.01
Humanity	0.9943	0.87	0.9543	0.12
Social Science	0.9898	0.98	0.838	0
STEM	0.9963	1	0.878	0

Table 7: Performance comparison on different datasets with and without adversarial training.

Dataset	Model	PPL↓	GRUEN↑	BS↑
Essay	Original	25.10	0.371	0.830
	Llama-3-8B	21.01	0.362	0.856
	Qwen2.5-7B	20.70	0.587	0.874
	GPT-5 (Ours)	18.90	0.661	0.900
Reuters	Original	8.42	0.678	0.904
	Llama-3-8B	9.09	0.679	0.902
	Qwen2.5-7B	9.09	0.679	0.902
	GPT-5 (Ours)	9.09	0.679	0.902
WP	Original	24.68	0.475	0.822
	Llama-3-8B	19.67	0.493	0.831
	Qwen2.5-7B	19.74	0.585	0.838
	GPT-5 (Ours)	20.80	0.647	0.897
Humanity	Original	30.81	0.365	0.819
	Llama-3-8B	19.81	0.564	0.851
	Qwen2.5-7B	21.46	0.631	0.862
	GPT-5 (Ours)	20.67	0.640	0.905
Social	Original	22.55	0.457	0.800
	Llama-3-8B	15.80	0.575	0.843
	Qwen2.5-7B	18.63	0.699	0.854
	GPT-5 (Ours)	17.54	0.729	0.899
STEM	Original	22.89	0.463	0.783
	Llama-3-8B	18.67	0.500	0.815
	Qwen2.5-7B	21.82	0.606	0.810
	GPT-5 (Ours)	20.08	0.643	0.819

Table 8: Ablation study on Stage 4 refinement models.

of MASH-generated texts at the cost of reduced accuracy on unattacked samples.

E Visualization of Rewritten Texts

Figure 13 shows the t-SNE projection of the feature representations from the RoBERTa detector. The plots reveal a clear distribution shift: original AI texts (red) are well separated from human texts (blue), whereas MASH-polished texts (green) successfully migrate into the human cluster, validating the effectiveness of our style transfer.

F Additional Ablation Studies

Alternative Refinement Models. Table 8 compares text quality across different refinement models to evaluate open-source alternatives for Stage 4.

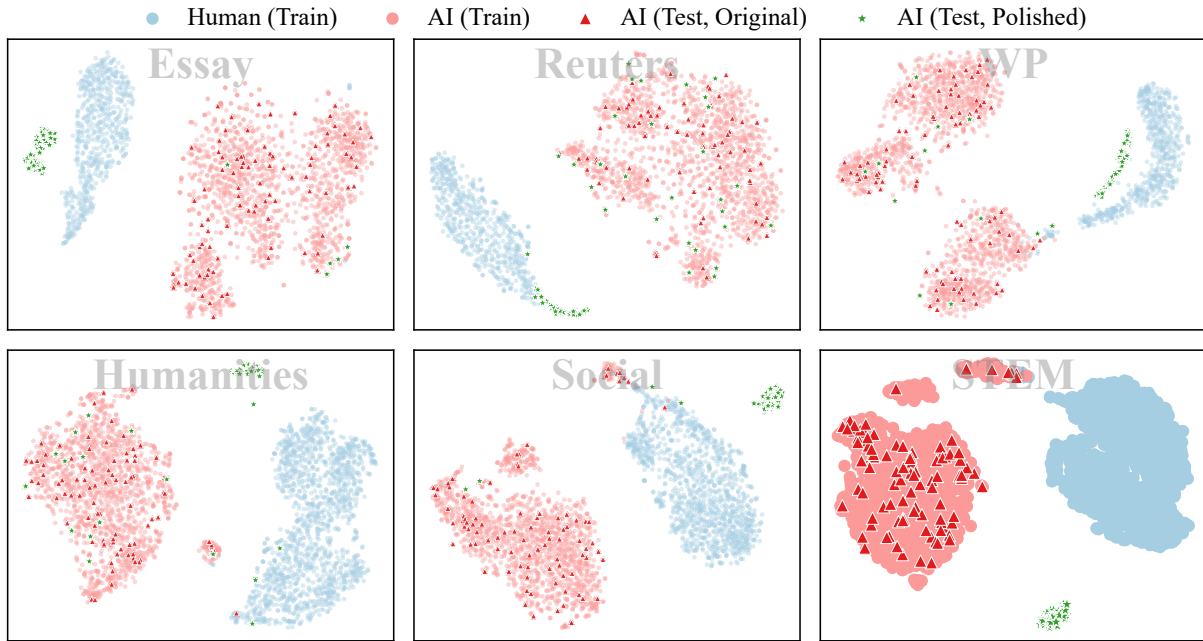


Figure 13: t-SNE visualization of feature representations extracted from the RoBERTa detector across six domains.

We assess Llama-3-8B-Instruct and Qwen2.5-7B-Instruct against our default GPT-5 refinement. Results show that both open-source models improve over original outputs, with Qwen2.5-7B achieving competitive GRUEN scores. However, GPT-5 consistently yields superior BERTScore, indicating better semantic preservation. These findings suggest that open-source refiners offer viable alternatives under cost constraints, though with a modest trade-off in semantic fidelity.

Impact of MASH Components. Table 9 evaluates the contributions of Style-SFT, DPO alignment, and inference-time refinement across three detectors and six datasets. The results demonstrate that the DPO module is critical for crossing detector decision boundaries, yielding a significant boost in ASR; however, this improvement often degrades textual fluency. The inference-time refinement module effectively mitigates this trade-off.

Generalizability of Refinement Module. Table 10 evaluates whether the inference-time refinement module (Stage 4) generalizes beyond MASH to other baseline attack methods. We apply the refinement module to outputs generated by ten baseline attacks across six datasets. The results demonstrate consistent improvements in text quality: PPL decreases substantially (e.g., TextFooler on Social: 136.89 \rightarrow 47.22), and GRUEN scores improve across most methods. BERTScore remains stable or improves slightly, indicating that semantic

consistency is preserved during refinement.

G Extended Analysis on Transferability

This section analyzes the transferability anomalies observed in Section 4.4. While MASH generalizes effectively across creative and argumentative domains (e.g., Essay), it fails to transfer to Reuters. We attribute this to stylistic mismatch: the rhetorical diversity learned by MASH conflicts with the rigid, objective, and concise standards of news reportage, rendering adversarial samples easily detectable. This mismatch also explains why generic corpora (dmitva, SemEval) failed to bypass Reuters-trained detectors, as they likely misclassify non-journalistic styles as AI-generated.

Regarding detector transferability (Figure 9), attacks transfer from supervised to zero-shot methods, but not vice versa. Supervised detectors rely on high-dimensional semantic features, compelling MASH to perform deep stylistic alignment. Since zero-shot detectors utilize simpler statistical proxies inherent to human writing, samples satisfying the rigorous semantic constraints of RoBERTa naturally meet the statistical thresholds of Binoculars.

H Real Cases

We randomly sample examples from our experiments, showing the original AI-generated texts alongside their MASH-polished versions to illustrate the stylistic transformation.

Detector	Method Variant	MGTBench-Essay				MGTBench-Reuters				MGTBench-WP			
		ASR	PPL	GRUEN	BS	ASR	PPL	GRUEN	BS	ASR	PPL	GRUEN	BS
RoBERTa	Style-SFT	0.83	25.10	0.5033	0.8823	0.09	8.58	0.6968	0.9930	0.68	19.79	0.5307	0.8781
	+ DPO	0.95	25.10	0.3711	0.8301	0.73	8.42	0.6779	0.9044	0.90	24.68	0.4748	0.8222
	+ Refine (Ours)	0.95	18.90	0.6614	0.9004	0.73	9.09	0.6790	0.9015	0.90	20.80	0.6471	0.8974
SCRN	Style-SFT	0.33	22.40	0.4615	0.8881	0.62	26.12	0.4912	0.8682	0.39	19.80	0.5727	0.9384
	+ DPO	0.56	23.11	0.4611	0.8679	0.83	26.26	0.4590	0.8496	0.81	20.34	0.5730	0.9258
	+ Refine (Ours)	0.56	13.30	0.6984	0.8951	0.83	14.01	0.6951	0.8964	0.81	12.70	0.6811	0.9433
Binoculars	Style-SFT	0.88	13.73	0.5143	0.8682	0.92	13.90	0.5293	0.8472	0.82	13.69	0.5889	0.8982
	+ DPO	0.94	13.24	0.4914	0.8540	0.95	13.32	0.4432	0.8335	0.85	13.58	0.5477	0.8809
	+ Refine (Ours)	0.94	11.39	0.6755	0.9077	0.95	12.91	0.7122	0.9037	0.85	12.05	0.6883	0.9286
Detector	Method Variant	MGT-Academic-Humanity				MGT-Academic-Social Science				MGT-Academic-STEM			
		ASR	PPL	GRUEN	BS	ASR	PPL	GRUEN	BS	ASR	PPL	GRUEN	BS
RoBERTa	Style-SFT	0.68	26.61	0.4366	0.8685	0.92	23.03	0.5738	0.8204	1.00	23.38	0.5221	0.7868
	+ DPO	0.87	30.81	0.3651	0.8193	0.98	22.55	0.4568	0.8004	1.00	22.89	0.4626	0.7833
	+ Refine (Ours)	0.87	20.67	0.6396	0.9048	0.98	17.54	0.7294	0.8993	1.00	20.08	0.6431	0.8194
SCRN	Style-SFT	0.70	31.77	0.5092	0.8955	0.65	27.24	0.4768	0.8816	0.75	38.57	0.4716	0.8721
	+ DPO	0.77	32.18	0.5072	0.8794	0.78	27.62	0.4742	0.8644	0.89	39.86	0.4374	0.8506
	+ Refine (Ours)	0.77	14.85	0.6762	0.9299	0.78	13.35	0.6626	0.9100	0.89	14.71	0.6488	0.9091
Binoculars	Style-SFT	0.92	15.16	0.5248	0.8454	0.95	16.02	0.5036	0.8480	0.94	17.60	0.4400	0.8280
	+ DPO	0.95	14.01	0.4184	0.8313	0.95	14.63	0.4445	0.8339	0.99	16.63	0.3508	0.8157
	+ Refine (Ours)	0.95	13.47	0.7025	0.9074	0.95	13.31	0.7014	0.9073	0.99	14.02	0.6434	0.8864

Table 9: Ablation study of MASH components across three detectors.

Method	MGTBench-Essay			MGTBench-Reuters			MGTBench-WP		
	PPL ↓	GRUEN ↑	BS ↑	PPL ↓	GRUEN ↑	BS ↑	PPL ↓	GRUEN ↑	BS ↑
DeepWordBug	63.02 → 62.57	0.470 → 0.485	0.934 → 0.933	88.83 → 91.32	0.524 → 0.524	0.933 → 0.929	32.85 → 30.41	0.524 → 0.553	0.950 → 0.951
TextBugger	40.67 → 33.32	0.457 → 0.516	0.919 → 0.932	84.02 → 88.84	0.404 → 0.404	0.893 → 0.890	34.24 → 24.15	0.494 → 0.532	0.932 → 0.945
TextFooler	77.18 → 72.24	0.447 → 0.480	0.940 → 0.943	181.46 → 197.37	0.302 → 0.303	0.870 → 0.868	50.20 → 37.01	0.425 → 0.490	0.938 → 0.946
DIPPER	44.78 → 44.76	0.648 → 0.651	0.903 → 0.906	11.07 → 11.07	0.718 → 0.718	0.910 → 0.910	14.11 → 14.09	0.577 → 0.583	0.901 → 0.902
Toblend	7.03 → 7.07	0.378 → 0.378	0.753 → 0.752	6.88 → 6.88	0.558 → 0.558	0.724 → 0.724	6.59 → 6.65	0.481 → 0.487	0.731 → 0.737
PromptAttack	11.21 → 11.99	0.778 → 0.772	0.916 → 0.915	10.55 → 11.41	0.791 → 0.791	0.927 → 0.924	10.55 → 11.27	0.726 → 0.726	0.925 → 0.921
Charmer	38.00 → 37.40	0.553 → 0.555	0.966 → 0.963	114.57 → 115.62	0.583 → 0.584	0.962 → 0.960	17.26 → 17.02	0.579 → 0.566	0.976 → 0.971
DPO-Evader	5.14 → 5.27	0.154 → 0.154	0.934 → 0.930	4.53 → 4.65	0.149 → 0.149	0.926 → 0.923	5.41 → 5.52	0.204 → 0.202	0.923 → 0.918
HMGC	57.80 → 49.17	0.588 → 0.604	0.923 → 0.930	82.60 → 88.63	0.547 → 0.547	0.901 → 0.898	41.82 → 37.87	0.576 → 0.573	0.932 → 0.931
GradEscape	17.72 → 16.76	0.135 → 0.191	0.904 → 0.911	74.96 → 84.20	0.481 → 0.482	0.814 → 0.815	13.49 → 13.59	0.652 → 0.648	0.987 → 0.981
CoPA	29.42 → 30.33	0.689 → 0.691	0.875 → 0.876	41.64 → 41.97	0.694 → 0.694	0.877 → 0.876	25.47 → 23.23	0.633 → 0.649	0.881 → 0.891
Method	MGT-Academic-Humanity			MGT-Academic-Social Science			MGT-Academic-STEM		
	PPL ↓	GRUEN ↑	BS ↑	PPL ↓	GRUEN ↑	BS ↑	PPL ↓	GRUEN ↑	BS ↑
DeepWordBug	30.09 → 24.03	0.487 → 0.517	0.952 → 0.959	32.59 → 22.18	0.541 → 0.600	0.951 → 0.963	22.64 → 19.88	0.466 → 0.489	0.954 → 0.958
TextBugger	40.51 → 22.91	0.403 → 0.507	0.932 → 0.958	56.87 → 26.43	0.443 → 0.570	0.936 → 0.959	39.93 → 27.79	0.404 → 0.479	0.933 → 0.946
TextFooler	72.33 → 31.04	0.350 → 0.494	0.929 → 0.955	136.89 → 47.22	0.332 → 0.525	0.912 → 0.952	81.24 → 45.71	0.319 → 0.448	0.926 → 0.946
DIPPER	14.67 → 14.38	0.578 → 0.595	0.903 → 0.925	13.88 → 13.31	0.683 → 0.673	0.906 → 0.926	12.47 → 11.66	0.604 → 0.646	0.879 → 0.901
PromptAttack	12.46 → 13.17	0.689 → 0.688	0.928 → 0.930	11.61 → 12.26	0.690 → 0.680	0.930 → 0.941	10.39 → 10.51	0.653 → 0.648	0.934 → 0.941
Charmer	14.54 → 13.65	0.550 → 0.545	0.971 → 0.967	12.88 → 11.83	0.641 → 0.599	0.974 → 0.966	11.62 → 11.57	0.544 → 0.530	0.987 → 0.983
DPO-Evader	9.82 → 9.91	0.441 → 0.441	0.921 → 0.921	6.80 → 7.12	0.353 → 0.366	0.927 → 0.928	8.37 → 8.97	0.540 → 0.534	0.933 → 0.934
HMGC	48.02 → 37.54	0.518 → 0.556	0.916 → 0.930	58.85 → 28.65	0.571 → 0.633	0.917 → 0.950	42.69 → 33.06	0.528 → 0.557	0.926 → 0.937
GradEscape	21.51 → 17.82	0.615 → 0.635	0.958 → 0.963	14.69 → 12.31	0.702 → 0.676	0.967 → 0.970	20.40 → 15.98	0.616 → 0.620	0.950 → 0.958
CoPA	27.52 → 21.76	0.639 → 0.667	0.873 → 0.904	29.95 → 17.80	0.635 → 0.696	0.874 → 0.923	23.75 → 18.40	0.645 → 0.681	0.870 → 0.901

Table 10: Performance comparison of baselines before and after applying the GPT-5 refinement model (Stage 4 of MASH). Metrics shown are **PPL** ↓ (lower is better), **GRUEN** ↑ (higher is better), and **BS** (BertScore) ↑ (higher is better). Format: *Original* → *Refined*.

Real cases of samples from MGTBench

Original Text: The CALM Act of 2012 was designed to regulate the volume of commercials on television. The Act limits the number of times a commercial can run and the length of the commercial. However, some advertisers have used alternative methods to circumvent the restrictions of the Act, including splicing a quick video at the end of a commercial break, using a device called a time-shift button or a software application that allows them to skip through commercials. During the COVID-19 pandemic, there has been an increase in complaints from viewers about the volume of sponsored content on television. Viewers are frustrated by the frequency and intensity of ads, which they view as disrespectful and invasive. In response, television networks have increased their efforts to create programs that are well-crafted and engaging, while also offering valuable information or entertainment value. The limitations of the CALM Act of 2012 include the fact that many viewers still find it difficult to understand how to change their cable or satellite settings to avoid seeing commercials altogether. Additionally, some advertisers are able to use loopholes or technical devices to circumvent the restrictions of the Act, while others simply continue to overuse commercials on a regular basis. A possible solution to address these limitations is for television networks to require advertisers to provide information on how

Polished Text: The CALM Act of 2012 was created to regulate the volume of commercials on television. The Act limits the number of times a commercial can run and its duration. However, some advertisers use alternative methods to bypass the restrictions of the CALM Act. For instance, advertisers may splice a quick video at the end of a commercial break or use devices like time-shift buttons or software applications to skip through commercials. During the COVID-19 pandemic, there has been an increase in complaints from viewers about the volume and frequency of sponsored content on television. Despite the CALM Act, many viewers struggle to understand how to adjust their cable or satellite settings to avoid commercials altogether. Others continue to overuse commercials, frustrating viewers further. A potential solution could involve television networks requiring advertisers to provide clear instructions on how viewers can manage or avoid commercials. This approach could help address some of the limitations of the CALM Act. While challenges remain, finding effective solutions to regulate commercials is crucial for improving the viewer experience.

Original Text: Ice Cream Industry Suffers Major Losses Due to Dull Summer in Europe, Prompting Big Players to Turn to Winter Marketing Strategies The ice cream industry in Europe is scrambling to recoup major losses from a lackluster summer season this year. Unseasonably cool and rainy weather across much of Northern and Western Europe led to a sharp drop in ice cream sales over the normally busy summer months. Sales of ice cream and related frozen desserts fell 12% in June and July compared to the same period last year, according to market research firm Kantar. The biggest players in the European ice cream market are turning to winter marketing campaigns to try to make up some of the shortfall. Unilever, owner of popular ice cream brands like Magnum, Ben & Jerry's, and Solero, reported its ice cream division's profits dropped 22% in the first half of the year. A spokesperson for Unilever said the company is relying more on promotional pricing, limited batch seasonal flavors, and partnerships with European cafes and bakeries to market ice cream during the fall and winter. Nestle, another dominant player that owns brands like Haagen-Dazs and Skinny Cow, said their ice cream profits fell 18% this summer. They are also pivoting to target ice cream sales in colder months. "We recognize that ice cream is not just a summer treat anymore," said Beatrice Guillaume-Grabisch, head of Nestle's global ice cream division. "Consumers want to indulge year-round, so we are creating more demand during winter by releasing popular seasonal flavors earlier, bundling smaller sizes at lower price points, and increasing our advertising in the winter months." Other companies are following suit with their own winter campaigns to boost falling ice cream revenues. Mars Wrigley, maker of brands like Snickers and Twix ice cream bars, is offering promotions like holiday gift boxes and limited Christmas flavors. Various local and craft ice cream makers across Europe are holding special "winter wonderland" sampling events and ice cream socials. A few brands are experimenting with warm-inspired flavors like spiced pumpkin, cinnamon, and nut mixes to resonate more with cold weather cravings. While winter marketing may help soften the blow, most industry experts expect significant long-term impacts from this summer's losses. "There's only so much these companies can do to counteract abysmal summer numbers," said Carlos Dominguez, an analyst specializing in the European frozen foods sector. "A dull summer in a normally hot weather country has massive consequences

that will likely continue hampering ice cream profits into next year and beyond." The ice cream industry is hoping for sunnier, warmer weather next summer to melt away the current cold front on their balance sheets.

Polished Text: Ice cream maker Nestle is reporting major losses due to a long period of winter in Europe. Big players in the ice cream industry are scrambling to recoup big losses from a lackluster summer season this year. Unseasonably cool and rainy weather across much of Northern and Western Europe led to a sharp drop in ice cream sales over the normally busy summer months. Sales of ice cream and related frozen desserts fell 12 percent in June and July compared with the same period last year, according to market research firm Kantar. The biggest names in the European ice cream market are turning to winter marketing campaigns to try to make up some of the shortfall. Nestle, owner of popular ice cream brands including Magnum, Ben & Jerry's and Solero, said its ice cream profits dropped 22 percent in the first half of this year alone. A spokesperson for Nestle said the company is relying more on promotional pricing, limited batch seasonal flavors and partnerships with European cafes and bakeries to market ice cream during the fall and winter. But Nestle, another dominant player in the market that owns brands including Haagen-Dazs and Skinny Cow, said their profits fell 18 percent this summer. It's also pivoting to target the dips in demand. "We recognize that ice cream is not just a summer treat anymore," said Beatrice Guillaume-Grabisch, head of Nestle's global ice cream division. "Consumers want to indulge year-round, so we are creating more demand during winter by releasing popular seasonal flavors earlier, bundling smaller sizes at lower price points and increasing our advertising in the winter months." Other companies are following suit with their own winter campaigns to boost profits. Mars Wrigley, maker of brands including Snickers and Twix ice cream bars, is offering promotions like holiday gift boxes and limited Christmas cards. Several local and craft ice cream makers across Europe are holding special "winter wonderland" sampling events and ice cream socials. A few brands are experimenting with warm-inspired flavors like spiced pumpkin, cinnamon and nut mixes to get the word out about cold weather cravings. While winter marketing may help soften the blow, most industry experts expect big impacts from this summer's record-setting fall.

Real cases of samples from MGT-Academic

Original Text: The death of Adolf Hitler and the looting of his headquarters have been subjects of considerable interest and controversy. One such case involves the Russian historian and author, Lev Bezymenski. In his works, including "Stalin and Hitler" (2002) and "The Death of Adolf Hitler: Unknown Documents from Soviet Archives" (1968), Bezymenski did not disclose that he had taken several containers of gramophone records from the Reich Chancellery. These recordings, which featured performances by renowned European and German orchestras and soloists, were among Hitler's favorite music pieces. Bezymenski, who was of Jewish descent, brought the looted collection back to Moscow. He reportedly felt guilty about his actions and hid the records in an attic. The collection remained hidden until it was accidentally discovered by his daughter, Alexandra Besymenskaja, in 1991. Bezymenski chose to keep silent about the records throughout his life, fearing accusations of looting. The collection is still in Russia. Another notable case of art looting during World War II involves Viktor Baldin, a Soviet army captain who later directed the Shchusev State Scientific Research Museum of Architecture in Moscow. On May 29, 1945, Baldin took 362 drawings and two small paintings from Karnzow Castle in Brandenburg. These artworks had been stored there by the Kunsthalle Bremen. Russian Culture Minister Mikhail Shvydkoi has estimated the value of the Baldin Collection to be around USD 1.5 billion. The looting of art and cultural artifacts during wartime is a complex and sensitive issue. It involves not only the moral and legal implications of such actions but also the challenges of restitution and the preservation of cultural heritage. The cases of Bezymenski and Baldin highlight the ongoing debates and efforts to address the legacy of art looting during World War II.

Polished Text: The death of Adolf Hitler and the looting of his headquarters have been subjects of considerable interest and controversy. Lev Bezymenski, a Russian historian and author, did not disclose in his works that he had taken several containers of gramophone records from the Reich Chancellery. These recordings featured performances by renowned European and German orchestras and soloists and were among Hitler's favorite music pieces. Bezymenski, who was of Jewish descent, brought the looted collection back to Moscow, where

he reportedly felt guilty about his actions and hid the records in an attic. The collection remained hidden until it was accidentally discovered by his daughter, Alexandra Besymenskaja, in 1991. Bezymenski chose to remain silent about the records throughout his life, fearing accusations of looting. The looted collection is still in Russia. Another case of art looting during World War II involves Viktor Baldin, a Soviet army captain who later directed the Shchusev State Scientific Research Museum of Architecture in Moscow. On May 29, 1945, Baldin took 362 drawings and two small paintings from Karnzow Castle in Brandenburg, which had been stored there by the Kunsthalle Bremen. The Baldin Collection, valued at approximately USD 1.5 billion according to Russian Culture Minister Mikhail Shvydkoi, highlights the immense cultural and financial significance of such looted artifacts. The looting of art and cultural artifacts during wartime raises complex moral, legal, and cultural heritage preservation issues, as seen in the cases of Bezymenski and Baldin.

Original Text: Interim Management: An Expert's Overview Interim management refers to the temporary provision of management resources and skills to an organization. It involves the appointment of experienced executives to manage specific transitions, projects, or periods of change within a company. Interim managers are typically hired for their expertise in areas such as: * Change management * Start-up and scale-up businesses * Sabbaticals * Management buyouts (MBOs) and Initial Public Offerings (IPOs) * Mergers and acquisitions * Project management The scope of an interim manager's skill set is unique, making their role invaluable in various business scenarios. The interim management concept has gained traction in several countries, including the UK, Germany, Belgium, Australia, the US, France, Ireland, Spain, and Nigeria. In Spain, the recruitment of interim managers increased by 68% in 2011, according to Michael Page Interim Management. Moreover, the Association Interim Management Spain was established in 2013, further solidifying the presence of interim management in the country. Similarly, the Institute Of Corporate and Interim Management in Nigeria, approved by The Federal Ministry Of Education, promotes Corporate and Interim Management practice in Africa. To foster international collaboration and networking among interim managers, Network INIMA - The International Network of Interim Manager Associations was launched in December 2020. This nonprofit organization unites leading European Interim Manager associations, representing nine countries and over 2500 practicing Interim Managers, many of whom operate internationally. Unlike other international networks created by providers, INIMA is run by Interim Managers for Interim Managers, adhering to common principles, values, and a code of conduct. In conclusion, interim management plays a crucial role in addressing various business challenges and transitions. With a growing presence in numerous countries, interim managers offer unique skills and expertise that contribute to organizational success. INIMA, as an International Network of Interim Manager Associations, aims to further strengthen the interim management community and its impact on global business practices.

Polished Text: Interim management refers to the temporary provision of management resources and skills to an organization. Interim managers are often hired to handle transitions, projects, or periods of change within a company. They bring expertise in areas such as change management, start-ups, scale-ups, and other specialized business needs. Organizations may hire interim managers to address specific challenges or to assist with management buyouts, IPOs, or mergers and acquisitions. Interim managers are also skilled in project management, making them valuable assets for various business scenarios. The demand for interim management has grown globally, with countries like the UK, Germany, Belgium, Australia, the US, France, Ireland, Spain, and Nigeria embracing the concept. In Spain, the recruitment of interim managers increased significantly, with a 68% rise reported in 2011 by Michael Page Interim Management. The establishment of the Association Interim Management Spain in 2013 further strengthened the presence of interim management in the country. In Nigeria, the Institute Of Corporate and Interim Management, approved by The Federal Ministry Of Education, promotes interim management practices across Africa. To foster international collaboration, Network INIMA was launched in December 2020 as a nonprofit organization. INIMA unites leading European Interim Manager associations, representing nine countries and over 2500 practicing Interim Managers. Unlike other networks created by providers, INIMA is managed by Interim Managers themselves and adheres to shared principles, values, and a code of conduct. Interim management plays a crucial role in addressing business challenges and transitions, contributing significantly to organizational success. INIMA aims to strengthen the interim management community and its impact on global business practices.