

LuxVeri at GenAI Detection Task 3: Cross-Domain Detection of AI-Generated Text Using Inverse Perplexity-Weighted Ensemble of Fine-Tuned Transformer Models

Md Kamrujjaman Mobin¹, Md Saiful Islam^{1,2}

¹Computer Science and Engineering, Shahjalal University of Science and Technology, Sylhet, Bangladesh

²Computing Science, University of Alberta, Edmonton, Alberta, Canada

Abstract

This paper presents our approach for Task 3 of the GenAI content detection workshop at COLING-2025, focusing on Cross-Domain Machine-Generated Text (MGT) Detection. We propose an ensemble of fine-tuned transformer models, enhanced by inverse perplexity weighting, to improve classification accuracy across diverse text domains. For Subtask A (Non-Adversarial MGT Detection), we combined a fine-tuned RoBERTa-base model with an OpenAI detector-integrated RoBERTa-base model, achieving an aggregate TPR score of 0.826, ranking 10th out of 23 detectors. In Subtask B (Adversarial MGT Detection), our fine-tuned RoBERTa-base model achieved a TPR score of 0.801, securing 8th out of 22 detectors. Our results demonstrate the effectiveness of inverse perplexity-based weighting for enhancing generalization and performance in both non-adversarial and adversarial MGT detection, highlighting the potential for transformer models in cross-domain AI-generated content detection.

1 Introduction

The proliferation of advanced language models such as GPT (Radford et al., 2019) and RoBERTa (Liu et al., 2019), machine-generated content has become prevalent across social media, journalism, and academia, raising concerns about authenticity and misinformation. Detecting AI-generated text is especially challenging across diverse domains, where variations in language and style can hinder detection efforts.

In Task 3 of the COLING 2025 Workshop on Detecting AI-Generated Content (Dugan et al., 2025), we tackle cross-domain Machine-Generated Text (MGT) detection using an ensemble approach that combines fine-tuned RoBERTa-base models (Liu et al., 2019) and OpenAI detection tools (Solaiman et al., 2019). Our method leverages inverse perplexity weighting to enhance the contributions of

high-confidence models, yielding a robust detection system.

Our approach achieved an aggregate score of 0.826 in Non-Adversarial Cross-Domain MGT detection (Subtask A), ranking 10th, and 0.801 in Adversarial Cross-Domain MGT detection (Subtask B), ranking 8th. This paper outlines our ensemble-based methodology, dataset considerations, and insights for effective cross-domain AI-generated text detection.

2 Background

2.1 Dataset

The RAID dataset (Dugan et al., 2024), provided for the competition, is designed for evaluating machine-generated text detectors. It contains over 10 million documents across 11 language models, 11 genres, 4 decoding strategies, and 12 adversarial attacks, including both human-written and machine-generated content from 8 different domains like books, news, poetry, and recipes. For training and validation, we used the RAID-train subset (802 million words, 11.8GB) and RAID-test subset (81 million words, 1.22GB). We also utilized the RAID-extra subset, which includes languages like Czech and German (275 million words, 3.71GB). This dataset provides a comprehensive resource for AI-generated text detection.

For the fine-tuning of our model, we reduced the dataset by using about 10% of the publicly available data. This reduction was carried out in a balanced manner across all genres, decoding strategies, attacks, and domains to ensure that each subset was proportionally represented. Specifically, we reduced the data across the following domains: abstracts, books, news, poetry, recipes, reddit, reviews, and wiki. The distribution of this reduced data across models is shown in Table 1, with domain-specific sample sizes for each model. For example, the number of samples for "Chat-

Model	Domain								Total
	Abstracts	Books	News	Poetry	Recipes	Reddit	Reviews	Wiki	
Human	2119	2137	2136	2125	2126	2135	1132	2135	17109
ChatGPT	4238	4274	4272	4250	4253	4270	2263	4270	34090
Cohere	4238	4274	4272	4250	4253	4270	2263	4270	34090
Cohere-Chat	4238	4274	4272	4250	4253	4270	2263	4270	34090
GPT-2	8477	8549	8544	8501	8506	8540	4526	8540	68183
GPT-3	4238	4274	4272	4250	4253	4270	2263	4270	34090
GPT-4	4238	4274	4272	4250	4253	4270	2263	4270	34090
Llama-Chat	8477	8549	8544	8501	8506	8540	4526	8540	68183
Mistral	8477	8549	8544	8501	8506	8540	4526	8540	68183
Mistral-Chat	8477	8549	8544	8501	8506	8540	4526	8540	68183
MPT	8477	8549	8544	8501	8506	8540	4526	8540	68183
MPT-Chat	8477	8549	8544	8501	8506	8540	4526	8540	68183
Total	62516	63577	63538	63380	63408	63606	37338	63606	518469

Table 1: Data distribution for various models across different domains, with total data per model and summed values for each domain. The values represent domain-specific sample sizes for each model. We used only 10% of the RAID (Dugan et al., 2024) dataset for fine-tuning our models.

GPT" in the "Books" domain is 4274, while for "Human" in the "Reviews" domain, it is 1132. This balanced reduction ensures the data used for training is representative across models and domains, enabling efficient and effective fine-tuning.

2.2 Related Work

The detection of machine-generated text has gained attention with the rise of large language models (LLMs) like GPT (Radford et al., 2019) and BERT (Devlin et al., 2019). Fine-tuned Transformer models have succeeded in binary classification tasks, but challenges remain in cross-domain and multilingual contexts due to data biases (Liu et al., 2019; Solaiman et al., 2019). Ensemble methods combining models like BERT, RoBERTa, GPT variants, and perplexity-based weighting have been explored to improve domain robustness (Schick and Schütze, 2020; Clark et al., 2019).

Recent work in cross-domain detection shows that RoBERTa-based detectors for GPT-2 generated technical text can be transferred with few labeled examples, such as from physics to biomedicine (Rodriguez et al., 2022). Paragraph-level detection is also being explored to address document tampering in mixed-domain texts.

For multilingual detection, models like XLM-RoBERTa (Conneau et al., 2019) and RemBERT (Chung et al., 2021) improve cross-lingual detection, though challenges remain for low-resource languages (Hu et al., 2020). Recent SemEval tasks (Fetahu et al., 2023; Wang et al., 2024) have refined these approaches with task-specific fine-tuning. Our work builds on these methods by using

inverse perplexity-weighted ensembles to enhance detection across domains and languages.

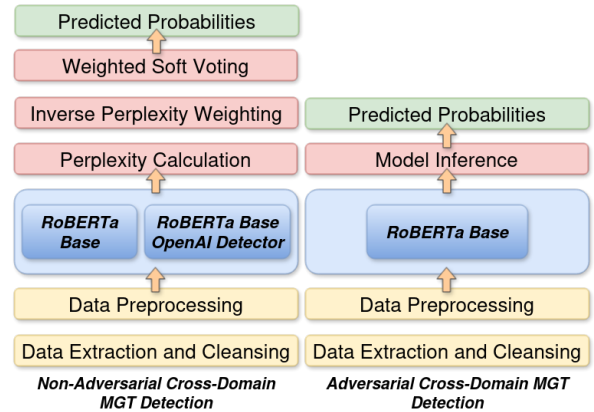


Figure 1: Overall Framework of our Proposed System for both Non-Adversarial and Adversarial Cross-Domain MGT Detection.

3 System Overview

We developed an ensemble approach for AI-generated text detection across multiple domains, using Transformer models with inverse perplexity-based weighted voting for improved accuracy. The system overview is shown in Figure 1.

3.1 Ensemble Model Selection and Justification

For the ensemble model, we selected two Transformer-based models tailored for Non-Adversarial and Adversarial cross-domain text detection, leveraging their strengths in capturing linguistic, syntactic, and semantic patterns essen-

tial for AI-generated content detection. For non-adversarial tasks, we employed **RoBERTa-base**, recognized for its strong performance in natural language understanding and robust generalization across text domains (Liu et al., 2019), alongside the **RoBERTa-base OpenAI Detector**, fine-tuned specifically for distinguishing AI-generated text from human-authored content (Solaiman et al., 2019). For adversarial scenarios, **RoBERTa-base** was selected again for its ability to capture subtle linguistic patterns and semantic inconsistencies, making it highly effective in challenging detection tasks (Liu et al., 2019). This strategic selection ensures a balanced approach to handling diverse and adversarial text detection challenges.

3.2 Data Pre-processing

For text classification, the data was preprocessed using model-specific tokenizers, incorporating truncation and padding as required. To enhance memory efficiency and training performance, texts were sorted by word count, reducing unnecessary padding. A fixed random seed was maintained to ensure reproducibility.

3.3 Training Procedure

The models were fine-tuned using the Hugging Face Transformers library¹ for English and multilingual text classification. Tokenization was performed with ‘AutoTokenizer’, and the architectures were adapted for classification tasks with appropriate label mappings.

Training was conducted for 3 epochs with a learning rate of 2×10^{-5} , batch sizes of 4 for training and 16 for validation, and weight decay of 0.01. Early stopping was applied with a patience of 5 evaluations and a 0.001 improvement threshold. Evaluation checkpoints were saved after each epoch, and the best-performing model was used for testing.

This procedure ensured robust generalization across subtasks. Further training details are provided in Table 2.

3.4 Ensemble Voting Strategy

Our ensemble employs a weighted soft-voting strategy, combining predictions from all fine-tuned models for each subtask. The weights are determined based on inverse perplexity, with lower perplexity values reflecting higher confidence.

¹Hugging Face Transformers: <https://huggingface.co/transformers/>

Hyperparameter	Value
Number of Epochs	2 ~ 3
Learning Rate	$1 \times 10^{-5} \sim 2 \times 10^{-5}$
Training Batch Size	4
Validation Batch Size	16
Early Stopping Patience	5 validation steps
Early Stopping Threshold	0.001
Weight Decay	0.01
Optimizer	AdamW
Loss Function	Binary Cross-Entropy
Evaluation Strategy	Every ¼ epoch
Checkpointing Strategy	Validation loss

Table 2: Training Configuration

3.4.1 Perplexity Calculation

For each model, we compute the perplexity based on its predictions. The perplexity P is computed using the Negative Log Likelihood formula:

$$P = \exp \left(-\frac{1}{N} \sum_{i=1}^N \log(p(y_i | x_i)) \right)$$

where $p(y_i | x_i)$ is the predicted probability for the true label y_i , and N is the number of test samples. Lower perplexity values indicate higher confidence.

To compute perplexity, we use each model’s logits, apply softmax to obtain probabilities, and then calculate perplexity based on the true labels and these probabilities.

3.4.2 Perplexity-Based Weighting Adjustment

To calculate model weights, each model’s perplexity is adjusted by subtracting 1, creating an effective weighting scale. The weight w_i for model i is then computed as the inverse of this adjusted perplexity and normalized across models, giving higher confidence models greater influence.

$$w_i = \frac{1/(P_i - 1)}{\sum_{j=1}^M (1/(P_j - 1))}$$

where M represents the total number of models, and P_i is the original perplexity of model i .

3.4.3 Weighted Soft-Voting

Each model’s predicted probabilities are scaled by its weight and summed to form the final ensemble prediction. This weighted voting prioritizes models with higher confidence (lower perplexity), giving them greater influence on the final decision. The ensemble’s final prediction for each class c is:

Non-Adversarial Results													
Detector	Chat GPT	GPT-4	GPT-3	GPT-2	Mistral	Mistral-Chat	Cohere	Cohere-Chat	Llama-Chat	MPT	MPT-Chat	AGG TPR	Rank
FT RoBERTa + RoBERTa OpenAI	0.960	0.861	0.895	0.753	0.734	0.936	0.546	0.748	0.891	0.804	0.901	0.826	10/23
FT RoBERTa + RoBERTa OpenAI + BERT	0.983	0.934	0.755	0.730	0.709	0.960	0.510	0.747	0.943	0.772	0.932	0.825	11/23
FT RoBERTa	0.943	0.836	0.902	0.739	0.719	0.916	0.542	0.737	0.870	0.800	0.891	0.813	12/23
Binoculars	0.997	0.907	0.989	0.678	0.610	0.914	0.935	0.943	0.973	0.447	0.707	0.790	-
Adversarial Results													
FT RoBERTa	0.911	0.808	0.873	0.730	0.720	0.887	0.567	0.740	0.855	0.806	0.861	0.801	8/22
FT RoBERTa + RoBERTa OpenAI	0.876	0.777	0.813	0.690	0.681	0.851	0.518	0.696	0.823	0.757	0.817	0.760	10/22
FT RoBERTa + RoBERTa OpenAI + BERT	0.896	0.843	0.675	0.663	0.651	0.874	0.457	0.670	0.857	0.711	0.841	0.749	11/22
SuperAnnotate AI Detector	0.963	0.913	0.720	0.411	0.342	0.897	0.445	0.685	0.918	0.314	0.767	0.649	-

Table 3: Cross-domain MGT detection performance under non-adversarial and adversarial conditions. The table shows detector performance across various generator models, with aggregate True Positive Rate (AGG TPR) and rankings. "FT" denotes fine-tuned models, and base models are used for training and evaluation.

$$p_{\text{ensemble}}(c) = \sum_{i=1}^M w_i \cdot p_i(c)$$

where $p_i(c)$ is the predicted probability for class c by model i , and w_i is its weight.

This method enhances ensemble accuracy by prioritizing predictions from more confident models, improving overall performance.

4 Results

Table 3 shows cross-domain MGT detection performance for non-adversarial and adversarial testing, with detectors ranked based on aggregate True Positive Rate (TPR).

4.1 Performance

In the non-adversarial setting, the fine-tuned RoBERTa + RoBERTa OpenAI model which was fine-tuned on RAID dataset (Dugan et al., 2024) achieved the highest performance, with an aggregate (AGG) score of 0.826, ranking 10th out of 23 detectors (see Table 3). This model effectively combined fine-tuned RoBERTa Base and RoBERTa Base OpenAI models, with perplexity-based weighting to give more influence to lower-

perplexity models, enhancing overall accuracy. It consistently delivered strong results across various generator models, including ChatGPT and GPT-3.

In the adversarial testing, the fine-tuned RoBERTa model outperformed other detectors, achieving an aggregate (AGG) score of 0.801 and ranking 8th out of 22 (see Table 3). This demonstrates the model’s robust adaptability in adversarial conditions, achieving top scores with GPT-3 and GPT-4, even under altered input scenarios.

4.2 Model Comparison

The performance of various detectors was evaluated under both non-adversarial and adversarial conditions, revealing key insights into their strengths and limitations.

In the non-adversarial setting, FT RoBERTa + RoBERTa OpenAI emerged as the top performer, achieving an AGG TPR of 0.826 and ranking 10th overall. It demonstrated exceptional performance with models such as ChatGPT (TPR: 0.960) and GPT-4 (TPR: 0.861), outperforming FT RoBERTa (AGG TPR: 0.813, ranked 12th) and the ensemble model FT RoBERTa + RoBERTa OpenAI + BERT (AGG TPR: 0.825, ranked 11th). Interestingly, Binoculars (Hans et al., 2024) showed strong

results with specific generators like GPT-3 (TPR: 0.989) and Llama-Chat (TPR: 0.973). However, its inconsistent performance with other generators, such as MPT (TPR: 0.447), limited its reliability. In contrast, FT RoBERTa + RoBERTa OpenAI demonstrated stable results across all generators, including strong performances with Mistral-Chat (TPR: 0.936) and Cohere-Chat (TPR: 0.748), underscoring its robustness and versatility.

Under adversarial conditions, FT RoBERTa proved to be the most robust model, achieving an AGG TPR of 0.801 and ranking 8th overall. It excelled with GPT-3 (TPR: 0.873) and MPT-Chat (TPR: 0.861), outperforming FT RoBERTa + RoBERTa OpenAI (AGG TPR: 0.760, ranked 10th) and FT RoBERTa + RoBERTa OpenAI + BERT (AGG TPR: 0.749, ranked 11th). In comparison, the SuperAnnotate AI Detector ([SuperAnnotate, 2024](#)) delivered competitive results with ChatGPT (TPR: 0.963), but its performance was inconsistent, particularly with GPT-2 (TPR: 0.411) and Mistral (TPR: 0.342). These results emphasize the variability of some detectors when faced with adversarial data, highlighting the consistent reliability of FT RoBERTa.

The consistent dominance of FT RoBERTa + RoBERTa OpenAI in non-adversarial settings and FT RoBERTa in adversarial conditions underscores the importance of tailoring architectures to specific scenarios. While models like Binoculars and SuperAnnotate excelled in isolated cases, their lack of stability across diverse generators reinforces the value of robust, well-balanced models like FT RoBERTa. These findings suggest that future efforts should focus on further optimizing architectures to enhance cross-domain robustness and adversarial detection capabilities.

5 Limitations

Our approach, while effective, has several limitations. Focusing on RoBERTa models for fine-tuning and ensemble weighting excluded alternatives like RemBERT ([Chung et al., 2021](#)) and XLM-RoBERTa ([Conneau et al., 2019](#)), which might better handle longer sequences, noisy data, and multi-label tasks.

Due to computational constraints, we trained on a subset of the RAID dataset, limiting the model’s ability to capture its full diversity. Training on the full dataset could greatly improve detection performance, especially for underrepresented domains.

Performance variability across generator models (e.g., GPT-4 vs. Mistral) and limited multilingual capabilities highlight the need for better cross-domain generalization and robust multilingual detection. While the ensemble approach enhanced generalization, it increased computational overhead, warranting exploration of more efficient strategies in future work.

6 Discussion and Conclusion

In this paper, we proposed an ensemble-based approach for cross-domain MGT detection, combining fine-tuned RoBERTa Base and RoBERTa Base OpenAI detectors with inverse perplexity weighting. Our method achieved competitive results, ranking 10th and 8th in non-adversarial and adversarial tasks, respectively, in Task 3 of the GenAI content detection workshop at COLING-2025. Inverse perplexity weighting improved generalization by prioritizing more confident models across diverse domains. For non-adversarial tasks, we explored an inverse perplexity-based ensemble approach. However, the detectors in this ensemble underperformed compared to the fine-tuned RoBERTa model, highlighting the value of fine-tuning on task-specific data and suggesting avenues for refining ensemble techniques.

Our results show that transformer-based models, particularly RoBERTa, are effective for non-adversarial and adversarial MGT detection. For non-adversarial detection (Subtask A), we achieved a score of 0.826, and for adversarial detection (Subtask B), we scored 0.801. However, cross-domain detection remains challenging, especially with varied generator models and multilingual data. Our system performed well with generators like ChatGPT and GPT-4 but struggled with others like Cohere and Mistral, indicating the difficulty of detecting diverse machine-generated content.

Due to limited computational resources, we trained on a subset of the available data. Despite this, our models performed well, demonstrating the potential of our approach even with partial data. This work lays the foundation for further progress in MGT detection, especially in adversarial and cross-lingual settings. Future research can focus on enhancing multilingual capabilities, incorporating more diverse language models, and exploring dynamic ensemble strategies to improve performance across domains and attack scenarios.

References

- Hyung Won Chung, Thibault Févry, Henry Tsai, Melvin Johnson, and Sebastian Ruder. 2021. [Rethinking embedding coupling in pre-trained language models](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Dan So. 2019. [Electra: Pre-training text encoders as discriminators rather than generators](#).
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Unsupervised cross-lingual representation learning at scale](#). *CoRR*, abs/1911.02116.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, et al. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of NAACL-HLT 2019*, pages 4171–4186.
- Liam Dugan, Alyssa Hwang, Filip Trhlfk, Andrew Zhu, Josh Magnus Ludan, Hainiu Xu, Daphne Ippolito, and Chris Callison-Burch. 2024. [RAID: A shared benchmark for robust evaluation of machine-generated text detectors](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12463–12492, Bangkok, Thailand. Association for Computational Linguistics.
- Liam Dugan, Andrew Zhu, Firoj Alam, Preslav Nakov, Marianna Apidianaki, and Callison-Burch Chris. 2025. [Genai content detection task 3: Cross-domain machine generated text detection challenge](#). In *Proceedings of the 1st Workshop on GenAI Content Detection (GenAIDetect)*, Abu Dhabi, UAE. International Conference on Computational Linguistics.
- Besnik Fetahu, Sudipta Kar, Zhiyu Chen, Oleg Rokhlenko, and Shervin Malmasi. 2023. [SemEval-2023 task 2: Fine-grained multilingual named entity recognition \(MultiCoNER 2\)](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 2247–2265, Toronto, Canada. Association for Computational Linguistics.
- 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](#). *Preprint*, arXiv:2401.12070.
- Xiang Hu, Xuezhe Liu, Peng He, and et al. 2020. [Xtreme: A massively multilingual evaluation benchmark](#).
- Y. Liu et al. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Alec Radford, Karthik Narasimhan, Tim Salimans, et al. 2019. [Language models are unsupervised multitask learners](#).
- Juan Diego Rodriguez, Todd Hay, David Gros, Zain Shamsi, and Ravi Srinivasan. 2022. [Cross-domain detection of GPT-2-generated technical text](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1213–1233, Seattle, United States. Association for Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2020. [Exploiting cloze-style pretraining for few-shot text classification](#). In *Proceedings of ACL 2020*, pages 255–265.
- Irene Solaiman, Miles Brundage, Jack Clark, Amanda Askell, Ariel Herbert-Voss, Jeff Wu, Alec Radford, Gretchen Krueger, Jong Wook Kim, Sarah Kreps, et al. 2019. [Release strategies and the social impacts of language models](#). *arXiv preprint arXiv:1908.09203*.
- SuperAnnotate. 2024. [Ai detector](#). Accessed: 2024-11-15.
- 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. 2024. [Semeval-2024 task 8: Multidomain, multimodel and multilingual machine-generated text detection](#). *Preprint*, arXiv:2404.14183.

A Appendix

Tools & Libraries	Version
Python	3.10.14
Pandas	2.2.2
NumPy	1.26.4
PyTorch	2.4.0
Transformers	4.44.2
Evaluate	0.4.3
WandB	0.16.6

Table 4: Main tools and libraries used in our system

Table 4 provide the details about the corresponding libraries, which are beneficial to help replicate our experiments.