



Explainable AI: XAI-guided context-aware data augmentation

Melkamu Abay Mersha^{a,*}, Mesay Gemeda Yigezu^b, Atnafu Lambebo Tonja^c,
Hassan Shakil^a, Samer Iskander^a, Olga Kolesnikova^b, Jugal Kalita^a

^a College of Engineering and Applied Science, University of Colorado Colorado Springs, Colorado Springs, 80918, CO, USA

^b Instituto Politécnico Nacional (IPN), Centro de Investigación en Computación (CIC), 07738, Mexico City, Mexico

^c Mohamed bin Zayed University of Artificial Intelligence (MBZUAI), Abu Dhabi, UAE

ARTICLE INFO

Keywords:

XAI
Explainable artificial intelligence
Interpretable
Deep learning
Machine learning
Neural networks
Data augmentation
Synonym replacement
Back translation
Large language models
LLMs
And natural language processing

ABSTRACT

Explainable AI (XAI) has emerged as a powerful tool for improving the performance of AI models, going beyond providing model transparency and interpretability. The scarcity of labeled data remains a fundamental challenge in developing robust and generalizable AI models, particularly for low-resource languages. Conventional data augmentation techniques introduce noise, cause semantic drift, disrupt contextual coherence, lack control, and lead to overfitting. To address these challenges, we propose XAI-Guided Context-Aware Data Augmentation. This novel framework leverages XAI techniques to modify less critical features while selectively preserving most task-relevant features. Our approach integrates an iterative feedback loop, which refines augmented data over multiple augmentation cycles based on explainability-driven insights and the model performance gain. Our experimental results demonstrate that XAI-SR-BT and XAI-PR-BT improve the accuracy of models on hate speech and sentiment analysis tasks by 6.6 % and 8.1 %, respectively, compared to the baseline, using the Amharic dataset with the XLM-R model. XAI-SR-BT and XAI-PR-BT outperform existing augmentation techniques by 4.8 % and 5 %, respectively, on the same dataset and model. Overall, XAI-SR-BT and XAI-PR-BT consistently outperform both baseline and conventional augmentation techniques across all tasks and models. This study provides a more controlled, interpretable, and context-aware solution to data augmentation, addressing critical limitations of existing augmentation techniques and offering a new paradigm shift for leveraging XAI techniques to enhance AI model training.

1. Introduction

The rapid advancement of large language models (LLMs), such as GPT (Radford et al., 2018) and BERT (Devlin et al., 2018), has transformed various domains, including safety-critical applications. Despite their impressive capabilities, these models operate as black boxes, raising concerns about transparency, trustworthiness, and interpretability. Explainable Artificial Intelligence (XAI) has emerged as a key solution to these concerns, offering insights into the decision-making processes of AI models. While XAI techniques have been successfully applied to various natural language processing (NLP) tasks, their full potential remains underexplored, particularly in the context of data augmentation (Bayer et al., 2023; Kwon & Lee, 2023). A major challenge in training robust AI models is the scarcity of annotated data, particularly for low-resource languages (Tonja et al., 2023). While high-resource languages such as English benefit from extensive, high-quality datasets, these advantages do not transfer easily to other linguistic groups. Consequently, AI models

often underperform in low-resource settings, exacerbating disparities in performance and limiting the global applicability of AI-driven solutions.

Data augmentation has been widely adopted to address data scarcity, improving model performance by artificially generating additional training data. Traditional augmentation techniques include synonym replacement, random insertion, random swap, and random deletion (Wei & Zou, 2019), which manipulate words or phrases to introduce variation. Back translation (Corbeil & Ghadivel, 2020; Sennrich, 2015) translates text into a high-resource language and back into the original language to generate paraphrased versions. Other approaches, such as data noising (Nishi et al., 2021), and sentence shuffling (Takahagi & Shinnou, 2023), introduce controlled perturbations to enhance generalization. More recently, LLMs have been employed to generate synthetic data through fine-tuning or prompting (Dai et al., 2025; Schick & Schütze, 2021). However, these methods face significant challenges in preserving linguistic and contextual integrity. LLM-based augmentation, while powerful, is prone to hallucinations—generating factually incorrect or

* Corresponding author.

E-mail addresses: mmersha@uccs.edu (M.A. Mersha), mgemedak2022@cic.ipn.mx (M.G. Yigezu), atnafu.tonja@mbzuai.ac.ae (A.L. Tonja), hshakil@uccs.edu (H. Shakil), siskande@uccs.edu (S. Iskander), kolesnikova@cic.ipn.mx (O. Kolesnikova), jkalita@uccs.edu (J. Kalita).

<https://doi.org/10.1016/j.eswa.2025.128364>

Received 4 March 2025; Received in revised form 25 May 2025; Accepted 26 May 2025

Available online 29 May 2025

0957-4174/© 2025 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

contextually irrelevant content—and struggles to maintain fidelity in low-resource languages due to underrepresentation in pre-training data (Mersha et al., 2024c; Wang et al., 2024). Conventional augmentation techniques, on the other hand, often introduce noise, disrupt semantic coherence, and lack adaptability, ultimately leading to overfitting or semantic drift. These challenges underscore the need for a more structured, explainability-driven approach to data augmentation.

To address these limitations, we introduce XAI-Guided Context-Aware Data Augmentation, a novel framework that leverages explainability techniques to guide the augmentation process. By identifying and preserving critical task-relevant features while modifying non-essential ones, our approach ensures semantic integrity and contextual coherence. Unlike conventional methods that apply augmentations indiscriminately, our method incorporates an iterative feedback loop that refines the augmented data based on both explainability insights and model performance gains. This structured augmentation process reduces noise, mitigates biases, and enhances the quality of synthetic data.

Furthermore, our technique facilitates cross-lingual transfer learning by ensuring semantic consistency across languages, making it particularly effective for low-resource NLP tasks. By maintaining the fidelity of linguistic structures, our method enables models trained on augmented data to generalize more effectively across multiple languages. Experimental evaluations demonstrate that our XAI-guided augmentation significantly improves model performance, outperforming conventional techniques while enhancing interpretability before and after augmentation. The key contributions of this study are as follows:

- We propose a novel **XAI-guided context-aware data augmentation technique**, pioneering the use of explainability to enable informed and context-driven augmentation strategies across languages.
- We introduce an **iterative refinement mechanism** that continuously improves augmentation quality through post-augmentation analysis.
- We mitigate noise, semantic drift, and overfitting by **preserving critical features** while modifying non-essential ones, ensuring model robustness and reducing biases.
- We enhance **cross-lingual transfer learning** by maintaining semantic fidelity across low-resource languages, improving generalization in multilingual AI applications.

The remainder of this paper is structured as follows: [Section 2](#) reviews the related work on data augmentation and XAI techniques. [Section 3](#) describes the methodology and experimental setup, covering datasets, models, and the selection of XAI techniques. [Section 4](#) presents the results, along with an evaluation of the proposed approach and a comparison with existing augmentation methods. [Section 5](#) outlines the limitations of the study and future research directions. Finally, [Section 6](#) concludes the paper.

2. Related work

We reviewed relevant studies on explainability techniques and existing data augmentation approaches. The explainability techniques subsection focuses on methods that provide insights into model predictions, which are crucial for developing XAI-guided context-aware data augmentation. The data augmentation approaches subsection provides an overview of existing augmentation methods, serving as a foundation for understanding how explainability can be leveraged to generate more effective augmented data.

2.1. Explainability techniques

XAI techniques address the black-box nature of AI models by revealing the internal workings and decision-making processes, providing insights into why and how models make predictions (Mersha et al.,

2024a,b). Feature-based XAI methods highlight key features influencing the model decisions, enhancing the interpretability of model outputs. Numerous approaches have been developed to generate explanations that indicate the relevance of each token in a model's prediction. For clarity and ease of understanding, we grouped these methods into four main categories: Perturbation, Activation, Gradient, and Attention (Fantozzi & Naldi, 2024; Mersha et al., 2025). Below, we would like to discuss a selection of prior works in this area briefly.

Perturbation-based. XAI methods alter input features and observe the impact on the model's output to assess their importance. These methods are model-agnostic and directly interpretable. SHAP (SHapley Additive exPlanations) uses game theory to attribute contributions of individual features to the output (Lundberg, 2017). LIME (Local Interpretable Model-Agnostic Explanations) perturbs input features and trains a surrogate linear model to approximate the original model locally around the input (Ribeiro et al., 2016). Occlusion Sensitivity masks parts of the input and observes the changes in output to identify critical regions (Zeiler, 2014).

Activation-based. XAI approaches leverage neural activations to trace the contribution of each input feature to the model predictions by analyzing neuron activations across the network. Class Activation Mapping (CAM) identifies input features that contribute most to model predictions by combining neuron activations with their associated weights (Zhou et al., 2016). Layer-wise Relevance Propagation (LRP) assigns relevance scores to input features by propagating the prediction score backward through the model (Bach et al., 2015). Concept Activation Vectors (CAVs) quantify the presence of human-interpretable concepts in the activations of neural network layers, enabling targeted explanations (Kim et al., 2018).

Gradient-based. XAI approaches leverage the gradients of a model's output with respect to its input to assess the contribution of input features to the model's decision-making process. Among these, Integrated Gradients (IG) computes the average gradient along a linear path from a baseline input to the actual input (Sundararajan et al., 2017). Gradient×Input enhances interpretability by multiplying the gradient of the output with respect to the input by the input values themselves, providing a direct measure of feature importance in the context of the model's predictions (Shrikumar et al., 2017). FullGrad extends traditional gradient methods by incorporating the gradient contributions of bias terms alongside input features, thereby offering a more comprehensive and detailed explanation (Srinivas & Fleuret, 2019). Building upon IG, Guided IG introduces layer-specific guided propagation, which enhances explanation clarity by refining gradient-based insights for deep networks (Kapishnikov et al., 2021).

Attention-based. XAI methods leverage the attention mechanisms within models to generate explanations, highlighting the parts of the input most relevant to the model's prediction. These methods are particularly well-suited for transformer architectures. Attention Flow and Attention Rollout leverage the attention mechanism in transformer models to explain their decisions. Attention Flow propagates attention scores through the layers to quantify the importance of tokens, while Attention Rollout accumulates attention weights across layers to trace the influence of input tokens on the output, providing a global explanation of the model's decision (Abnar & Zuidema, 2020). Attention mechanism visualization helps interpret how attention layers focus on input tokens (Vig, 2019).

Several hybrid approaches have been proposed by combining different techniques and leveraging their complementary strengths. These include Gradient + Activation (Selvaraju et al., 2020), Activation + Perturbation (Shrikumar et al., 2017), Activation + Attention (Chefer et al., 2021), Attention + Gradient (Qiang et al., 2022), Gradient + Perturbation (Sundararajan et al., 2017), and Attention + Gradient + Perturbation (Yuan et al., 2021). By integrating these

methods, researchers aim to address the limitations of individual techniques.

2.2. Data augmentation approaches

Text data augmentation is a crucial technique in natural language processing (NLP) that enhances model performance by artificially expanding training datasets (Khan & Venugopal, 2024). By generating diverse variations of text data, augmentation improves generalization, robustness, and efficiency, particularly in low-resource scenarios (Shorten et al., 2021). The augmentation techniques create diverse training samples through strategic transformations. As a result, they strengthen model robustness and effectiveness in real-world applications.

Several text augmentation techniques exist, each offering unique benefits. In this work, we present the most common and widely used augmentation approaches.

Synonym replacement. This method replaces words in a sentence with their synonyms to generate new variations (Bayer et al., 2022). Wei and Zou (2019) introduce the Easy Data Augmentation (EDA) technique, a straightforward method that uses random synonym replacement, random insertion, random swap, and random deletion. EDA enhances model performance, particularly on small datasets, by reducing overfitting and improving generalization. This approach has demonstrated significant accuracy improvements for CNN and RNN models.

Back translation. Back translation is a widely used data augmentation technique in NLP that translates text into another language and then back into the original language (Hayashi et al., 2018; Taheri et al., 2024). This process generates paraphrased text. Sennrich (2015) enhance neural machine translation by generating synthetic parallel data from monolingual corpora. It translates target-language text into the source language and then back to the target, improving model performance, especially in low-resource settings (Sennrich, 2015). Corbeil and Ghadivel (2020) introduce a simple back-translation technique for generating diverse paraphrased data to improve transformer-based models. The approach enhances paraphrase identification by leveraging back-translation to create augmented training data, which helps improve model generalization, particularly in low-resource scenarios (Corbeil & Ghadivel, 2020; Latief et al., 2024).

Contextual word embedding-based augmentation. This approach utilizes pre-trained language models, such as BERT, GPT, and T5 (Shakil et al., 2024), to generate augmented text by replacing words based on contextual embeddings. By considering the surrounding words, it ensures that the augmented text remains grammatically and semantically coherent (Kapusta et al., 2024). Kobayashi (2018) introduces Contextual Augmentation, a data augmentation technique that enhances text diversity by replacing words with contextually appropriate alternatives using language models. This method leverages paradigmatic word relations to generate meaningful variations. Contextual Augmentation improves model robustness and generalization by preserving sentence coherence, particularly in text classification tasks (Kobayashi, 2018). Kumar et al. (2020) explored pre-trained transformer models for data augmentation in low-resource NLP tasks, comparing BERT, GPT-2, and BART. Their approach prepends class labels to input text, enabling label-aware augmentation. BERT replaces masked tokens while preserving meaning, GPT-2 generates fluent text but struggles with label retention, and BART applies denoising autoencoding (word masking), achieving the best balance between semantic fidelity and data diversity (Kumar et al., 2020).

Adversarial example generation. This approach involves generating and incorporating adversarially perturbed text samples. Volpi et al. (2018) propose an adversarial data augmentation approach to improve model generalization to unseen domains without access to target domain data.

Their method iteratively augments training data by generating adversarial examples from fictitious “hard” target distributions, formulated using a worst-case optimization over distributions near the source domain in the feature space (Volpi et al., 2018). Alzantot et al. (2018) propose a genetic algorithm-based approach for generating adversarial examples in NLP by strategically modifying input text while preserving its semantic meaning. Their method involves synonym-based word substitutions guided by word embeddings, ensuring that alterations remain fluent and natural (Alzantot et al., 2018).

Sentence shuffling. This approach randomly changing the order of sentences within a document while maintaining coherence (Iyyer et al., 2015). Takahagi and Shinnou (2023) rearranges phrases in Japanese sentences while maintaining dependency relationships to enhance textual entailment tasks. This method was applied to the JSICK dataset using BERT (Devlin et al., 2018) and RoBERTa (Liu, 2019) models, improving performance, particularly in detecting contradictions.

Existing data augmentation techniques face several key challenges. Random synonym replacement, random insertion, swapping, and deletion often distort context and introduce noise, negatively impacting the model’s ability to understand semantic coherence (Ding et al., 2024; Nair et al., 2024). Back-translation can be expensive and yield incorrect outputs, and mixups tend to create unrealistic or implausible examples. Adversarial examples may introduce synthetic noise that does not align with real-world language patterns. LLM-based approaches also struggle to maintain contextual fidelity, avoid hallucinations, and manage biases (Hadi et al., 2024). Generally, these existing methods are often “blind” as they fail to preserve critical features essential for models’ decision-making.

These limitations highlight the need for our XAI-guided context-aware data augmentation technique, which leverages explainability methods to identify and modify *less critical features* in the model’s predictions. We generate diverse data without altering critical features by focusing on *non-essential features* for augmentation-using synonym replacement, back-translation, or paraphrasing. This approach minimizes noise, avoids unnatural examples, and improves model generalization, effectively balancing data diversity and semantic integrity.

3. Methodology

Fig. 1 provides a high-level overview of our XAI-guided context-aware data augmentation methodology. The augmentation process starts by leveraging pre-trained multilingual transformer-based encoder models as the baseline model. To better understand the model’s decision-making, XAI techniques are integrated to identify and highlight the less important features in the dataset. Once these less influential words are identified, the data augmentation process begins, involving multiple steps.

In the XAI Synonym Replacement with Back Translation (XAI-SR-BT) approach, the less influential words identified by the XAI method are first translated into a high-resource language (English) using the Google Translate API. For each translated word, contextually appropriate synonyms are retrieved using the WordNet lexical database via the NLTK interface (Bird et al., 2009; Fellbaum, 1998). These synonyms are then back-translated into the original language using the same translation API. Finally, the original less important words in the input text are replaced with the back-translated synonyms. This method introduces controlled lexical variation while preserving semantic equivalence.

In the XAI Paraphrasing Replacement with Back Translation (XAI-PR-BT) approach, the translation process follows the same steps as in XAI-SR-BT. The fine-tuned version of the pre-trained PEGASUS model (tuner007/pegasus_paraphrase) is then applied to generate paraphrases for the translated less influential words (Zhang et al., 2020). However, other pre-trained models, such as GPT, can also be used for phrase generation. PEGASUS is better suited for this task because of its architecture. Unlike GPT, which is designed for next token prediction and

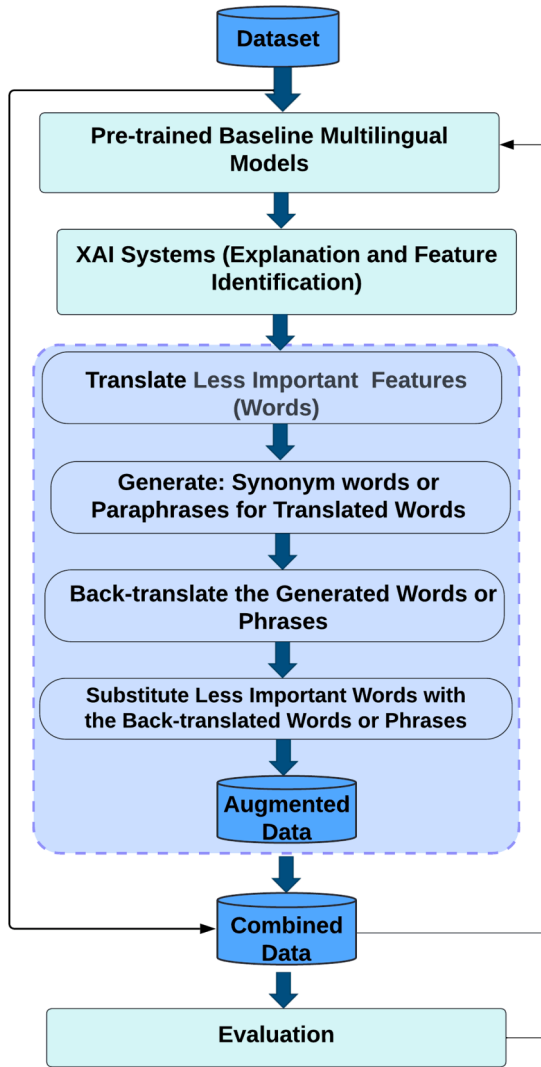


Fig. 1. XAI-guided and controlled context-aware data augmentation framework: An iterative XAI-guided feature identification data augmentation approach.

open-ended generation, PEGASUS is pre-trained with a gap-sentence generation objective, making it more effective in producing semantically consistent and contextually relevant phrases (Zhang et al., 2020). These paraphrased outputs are back-translated into the original language using the same translation API. Finally, the original less important words in the input text are replaced with the back-translated paraphrases. This technique enables more fluent and context-aware augmentation, thereby improving the model's robustness and generalization across linguistic variations.

The effectiveness of these augmentation strategies is evaluated by assessing the models' performance metrics and the explanations provided by the XAI techniques. Our approach incorporates an iterative feedback refinement loop into the XAI-guided augmentation pipeline, as illustrated in Fig. 1, to ensure the effectiveness and quality of augmentation. After performing the augmentation process, which is carried out through XAI-SR-BT and XAI-PR-BT, we dynamically verify whether the replacements for the selected less important words were successfully applied. If replacements are missing due to the unavailability of appropriate synonyms or paraphrases, the process loops back and dynamically adjusts the selection threshold for less important words identified by the XAI method (e.g., increasing the threshold from 20 % to 30 %).

Table 1

An overview of hate speech and sentiment analysis dataset statistics for multiple languages used in this study, including dataset size, average sentence length, and data sources.

Languages	Size	Average length	Data source
Hate speech dataset			
Amharic	14,643	17.5	Ayele et al. (2023)
Arabic	3353	13.6	Ousidhoum et al. (2019)
Hindi	8192	25.6	Bhardwaj et al. (2020)
Italian	8100	24.7	Bosco et al. (2018)
Portuguese	3059	15.7	Fortuna et al. (2019)
Spanish	6600	21.3	Basile et al. (2019)
Sentiment analysis dataset			
Amharic (Ethiopia)	9480	19.7	Muhammad et al. (2023)
Hausa (Nigeria)	22,153	18	Muhammad et al. (2023)
Kinyarwanda (Rwanda)	5155	14.6	Muhammad et al. (2023)
Swahili (Kenya)	3011	14	Muhammad et al. (2023)

Secondly, once the augmented data is generated, it is combined with the original dataset and used to retrain the model. The model is then evaluated using the accuracy and the F1 score metrics, along with explanation insights to assess the quality and interpretability of the predictions. If the evaluation results indicate insufficient improvement, the process enters a refinement phase again. In this iterative loop stage, we adjust the translation API, the paraphrasing model, or the XAI technique. This iterative process continues until the model demonstrates meaningful improvements. This flexibility is a key benefit of our methodology.

3.1. Datasets

We utilized six hate speech datasets spanning multiple languages, ranging from high-resource Spanish to low-resource Amharic, including Arabic, Hindi, Italian, and Portuguese. Additionally, to demonstrate the reproducibility of our approach, we employed four sentiment analysis datasets in Amharic, Hausa, Kinyarwanda, and Swahili. These datasets are carefully selected to evaluate the effectiveness of XAI-guided context-aware data augmentation techniques in high, medium, and low-resource languages. Each dataset differs in size, linguistic structure, and complexity, providing a robust and comprehensive benchmark for assessing the impact of our approach on model performance. This diversity enables us to investigate how well our proposed method generalizes across languages and varying levels of resource availability. Table 1 presents an overview of the characteristics and statistics of each dataset used in our experiments.

3.2. Models

Our model selection strategy is driven by two primary objectives: optimizing the data augmentation process and thoroughly evaluating the impact of the augmented data on model performance. To guide the augmentation process, we selected XLM-Roberta (XLM-R) (Conneau et al., 2019) as our baseline model due to its proven robustness and exceptional performance in cross-lingual tasks. We also considered other pre-trained multilingual models as they provide a strong foundation for language understanding and can significantly enhance performance in various language scenarios. We fine-tune these models on both the original and augmented datasets to adapt them to specific downstream tasks. This allows us to evaluate the impact of data augmentation on task-specific performance.

To evaluate the impact of our augmented data, we utilized XLM-R and mBERT (Pires et al., 2019), allowing for a comprehensive assessment of the effectiveness of our augmentation strategies across different architectures and languages. This dual-model evaluation enabled us to systematically analyze improvements in generalization, accuracy, and robustness across diverse languages, offering more profound insights

into how the augmented data influences model performance in different language scenarios. By employing a multi-model evaluation approach, we validate our augmentation's contribution to enhancing multilingual model performance and better understand the benefits of XAI-guided methodologies in different language settings.

3.3. Selection of explainability techniques

The selection of the appropriate XAI method is crucial and must be based on factors such as the interpretability needs, model architecture, and input data type (Yigezu et al., 2024). Several categories of XAI techniques provide distinct ways to interpret complex models, including perturbation-based methods like SHAP (Ribeiro et al., 2016), gradient-based approaches such as Integrated Gradients (IG) (Sundararajan et al., 2017), Layer-wise Relevance Propagation (Bach et al., 2015), and attention mechanism visualizations (Vaswani et al., 2017).

IG provides more accurate and interpretable word-level feature attributions than perturbation-based approaches, such as SHAP and LIME (Zhuo and Ge, 2024). IG is ideal for text-based models because it captures the contextual influence of each word on the model's predictions (Ansari et al., 2024). IG is well-suited for handling long sequential text inputs, as it leverages the full context of the input to compute attribution scores without the scalability limitations associated with perturbation-based methods such as SHAP and LIME (Sundararajan et al., 2017). Unlike SHAP and LIME, which rely on repeated input perturbations and extensive sampling to estimate feature importance (Bhattacharya, 2022), IG performs a single pass using gradient computations along a defined path from baseline to the actual input (Sundararajan et al., 2017). This makes IG significantly more computationally efficient and scalable for large-scale NLP tasks involving transformer-based models and long-form text data (Atanasova, 2024).

IG ensures that the selected features for augmentation have minimal impact on the model's overall decision-making process. IG is less sensitive to noise than perturbation-based methods like SHAP and LIME (Sundararajan et al., 2017), which can produce inconsistent feature importance scores due to their reliance on random perturbations. Perturbation-based XAI methods are also computationally expensive and often struggle with longer text sequences, leading to scalability issues and impractical for large-scale text data (Holzinger et al., 2022; Lundberg, 2017). Attention-based visualization techniques highlight words based on the attention weights assigned to each token by the model. However, these weights do not necessarily correlate with the true importance of features for the model's final prediction, which makes them unreliable for guiding XAI-guided context-aware data augmentation (Serrano & Smith, 2019). For instance, sentiment-bearing words like "excellent" and "poor" are crucial for determining classification in a classification task. In contrast, structural words like "had" or "but" might receive higher attention scores due to their contextual role, even though they contribute minimally to the text classification. Relying solely on attention scores for feature selection could lead to the preservation of irrelevant words and result in suboptimal data augmentation strategies.

We further contextualize our selection of IG by referencing our prior study (Mersha et al., 2025), in which we systematically evaluated the effectiveness of various XAI techniques across encoder-based language models. Our findings demonstrated that gradient-based techniques outperform SHAP, LIME, and attention visualization methods in terms of scalability and robustness in both short- and long-text inputs, particularly as the size of the model parameter increases (Mersha et al., 2025). In contrast, Integrated Gradients accurately quantify the contribution of each word, distinguishing between critical and non-essential features (Sundararajan et al., 2017). This capability allows us to modify or replace less important words while preserving the core meaning of the text. Thus, we determine that IG is the optimal choice for XAI-guided context-aware data augmentation due to its interpretability, computational efficiency, and ability to maintain the integrity of the augmented data (Ismail et al., 2021).

3.4. XAI-guided data augmentation strategy

The augmentation process begins using pre-trained multilingual models as the baseline model. We apply feature-based XAI techniques, such as Integrated Gradients, to explain the model's decision-making process and identify the features in the data that have minimal impact on its predictions. Building on this, we proposed two XAI-guided context-aware data augmentation methods: XAI Synonym Replacement with Back Translation and Paraphrasing Replacement with Back Translation. These methods leverage explainability to selectively augment the data by targeting features that contribute the least to the model's predictions. In both approaches, we employed the Integrated Gradients explainability technique to compute the importance score $I(f_i)$ for each feature (word) f_i within our dataset D_{original} . Integrated Gradients measure each feature's contribution to the model's prediction. To select the least important top- k features from a total of K features in the input example, we assumed that each feature has an importance score $I(f_i)$, where f_i is the i th feature, and $I(f_i)$ represents its importance. The features were ranked in ascending order based on their importance scores, resulting in an ordered list $f_{(1)}, f_{(2)}, \dots, f_{(K)}$, where $I(f_{(1)}) \leq I(f_{(2)}) \leq \dots \leq I(f_{(K)})$. From this list, we selected the top- k least important features for targeted data augmentation, denoted as $S_k = \{f_{(1)}, f_{(2)}, \dots, f_{(k)}\}$. The value of k is determined by input length, with shorter texts having smaller k values and longer texts having larger k values.

In the XAI Synonym with Back Translation augmentation approach, we replaced synonyms through translation after identifying the least important words. For each word $f_i \in S_k$, we translated it from the source language to English, represented as $f_{\text{trans}} = \text{Translate}(f_i)$. We then replaced the translated word with its synonym, denoted as $f_{\text{syn}} = \text{Synonym}(f_{\text{trans}})$, and back-translated the synonym back to source language as $f_{\text{bt}} = \text{BackTranslate}(f_{\text{syn}})$. The original least important word f_i in each sample $x \in D_{\text{original}}$ was replaced with its back-translated synonym f_{bt} .

We then generated augmented samples by replacing the least important words S_k in each original sample x with the back-translated synonyms, represented as $x_{\text{augmented}} = \text{Augment}(x, S_k)$. The augmented samples were collected into an augmented dataset, $D_{\text{augmented}}$.

In the final step, we combined the augmented dataset $D_{\text{augmented}}$ with the original dataset D_{original} to create the final dataset for further training and evaluation. This combined dataset is represented as $D_{\text{combined}} = D_{\text{original}} \cup D_{\text{augmented}}$.

The XAI Paraphrasing with Back Translation method follows a similar process to the synonym replacement approach, but uses paraphrasing instead. After identifying the top- k least important features S_k using Integrated Gradients, each word $f_i \in S_k$ was first translated from the source language into English ($f_{\text{trans}} = \text{Translate}(f_i)$). Instead of finding a synonym, we generated a paraphrase using pre-trained models, such as GPT or PEGASUS ($f_{\text{para}} = \text{Paraphrase}(f_{\text{trans}})$). The paraphrased word was then back-translated into the original source language, resulting in $f_{\text{bt}} = \text{BackTranslate}(f_{\text{para}})$. Each original word f_i was replaced with its back-translated paraphrased version f_{bt} , creating an augmented sample $x_{\text{augmented}} = \text{Augment}(x, S_k)$ for each original sample $x \in D_{\text{original}}$. As with the synonym replacement method, the augmented dataset $D_{\text{augmented}}$ can also be combined with the original dataset to form the final dataset $D_{\text{combined}} = D_{\text{original}} \cup D_{\text{augmented}}$.

Through these approaches, we effectively generated augmented data that leverages XAI techniques to target the least important features (words). While any features can also be targeted, our approach focuses on the least important words, preserving essential ones to generate model-oriented augmented data. This ensures the core information is retained while enhancing diversity and improving model robustness and performance. The model learns to be less sensitive to noise and irrelevant inputs by altering less critical features.

Table 2

Impact of XAI-guided synonym replacement and paraphrasing with back translation on XLM-R and mBERT models across different languages for the **hate speech detection task**. This table shows the improvements in Accuracy (Acc) and F1 score (F1) before and after augmentation, along with the corresponding delta values (ΔAcc and ΔF1).

Model	Dataset	Before augmentation		XAI synonym with back translation				XAI paraphrasing with back translation			
		Acc	F1	Acc	F1	ΔAcc	ΔF1	Acc	F1	ΔAcc	ΔF1
XLM-R	Amharic	0.865	0.864	0.931	0.931	+0.066	+0.067	0.911	0.911	+0.046	+0.047
	Arabic	0.775	0.769	0.839	0.836	+0.064	+0.067	0.814	0.811	+0.039	+0.042
	Hindi	0.812	0.810	0.887	0.885	+0.075	+0.075	0.863	0.860	+0.051	+0.050
	Italian	0.689	0.678	0.749	0.748	+0.060	+0.070	0.733	0.733	+0.044	+0.055
	Portuguese	0.712	0.709	0.796	0.795	+0.084	+0.086	0.766	0.761	+0.054	+0.052
	Spanish	0.735	0.735	0.809	0.809	+0.074	+0.074	0.773	0.771	+0.038	+0.036
mBERT	Amharic	0.870	0.872	0.925	0.925	+0.055	+0.053	0.913	0.912	+0.043	+0.040
	Arabic	0.779	0.766	0.824	0.823	+0.045	+0.057	0.814	0.812	+0.035	+0.046
	Hindi	0.845	0.834	0.897	0.892	+0.052	+0.058	0.881	0.880	+0.036	+0.046
	Italian	0.695	0.695	0.762	0.762	+0.067	+0.067	0.733	0.733	+0.038	+0.038
	Portuguese	0.675	0.669	0.744	0.742	+0.069	+0.073	0.721	0.711	+0.046	+0.042
	Spanish	0.765	0.765	0.827	0.822	+0.062	+0.057	0.803	0.795	+0.038	+0.030

Table 3

Impact of XAI-guided synonym replacement and paraphrasing with back translation on XLM-R and mBERT models across four low-resource languages for the **sentiment analysis task**. The table shows the Accuracy (Acc) and F1 score (F1) before and after augmentation, along with the corresponding delta values (ΔAcc and ΔF1).

Model	Dataset	Before augmentation		XAI synonym with back translation				XAI paraphrasing with back translation			
		Acc	F1	Acc	F1	ΔAcc	ΔF1	Acc	F1	ΔAcc	ΔF1
XLM-R	Amharic	0.803	0.801	0.884	0.883	0.081	0.082	0.863	0.854	0.060	0.053
	Hausa	0.827	0.825	0.907	0.896	0.080	0.071	0.875	0.874	0.048	0.049
	Kinyarwanda	0.706	0.700	0.764	0.761	0.058	0.061	0.749	0.741	0.043	0.041
	Swahili	0.648	0.626	0.687	0.679	0.039	0.053	0.671	0.663	0.023	0.037
mBERT	Amharic	0.814	0.811	0.892	0.891	0.078	0.080	0.872	0.870	0.058	0.059
	Hausa	0.819	0.814	0.914	0.913	0.095	0.099	0.884	0.881	0.065	0.067
	Kinyarwanda	0.717	0.710	0.776	0.774	0.059	0.064	0.751	0.743	0.034	0.033
	Swahili	0.661	0.653	0.698	0.693	0.037	0.040	0.689	0.686	0.028	0.033

4. Experimental results and discussion

We conducted two experiments to generate augmented data using XAI-guided context-aware data augmentation techniques with the XLM-R model across six languages. The first experiment applied XAI-guided Synonym Replacement with Back Translation, while the second employed XAI-guided Paraphrasing with Back Translation.

We measured improvements in accuracy and F1 score using both XLM-R and mBERT models to evaluate the effectiveness of the XAI-guided augmentation techniques. The results from both experiments, presented in Tables 2 and 3, demonstrate that models trained on XAI-guided augmented data consistently outperformed models trained on non-augmented data. The XAI Synonym Replacement with Back Translation approach achieved the highest performance gains across all languages for both XLM-R and mBERT models, showing larger improvements in accuracy and F1 score compared to the XAI Paraphrasing with Back Translation technique. The XAI Synonym Replacement with the Back Translation approach captures semantic variations more effectively, leading to better generalization and stronger model performance. Although XAI Paraphrasing with Back Translation also produced positive results, its impact was generally smaller than XAI Synonym Replacement with the Back Translation approach, which may not add sufficient diversity to the augmented data, as shown in Fig. 2.

These performance gains can be attributed to XAI-guided context-aware data augmentations' more informed and targeted nature. Unlike conventional methods that apply augmentations without accounting for feature (word) relevance for the decision-making process, which can introduce noise and distort the augmented data, our XAI-guided techniques leverage explainability methods, such as integrated gradients, to focus on altering less significant words in the input text. This selective approach ensures that the most critical features (words) are preserved

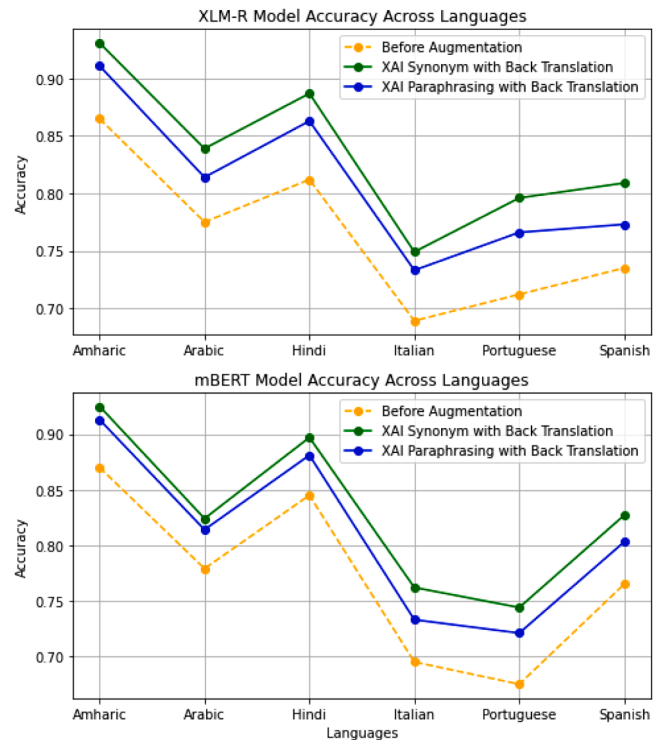


Fig. 2. Accuracy comparison of XLM-R and mBERT models across multiple languages on a hate speech dataset, before and after XAI-guided context-aware data augmentation, showcasing the impact of augmented data on model performance.

Table 4

Performance comparison of conventional and XAI-guided context-aware data augmentation techniques on XLM-R and mBERT models across different languages for the **hate speech detection task**. The table presents the accuracy (Acc) and F1 score (F1), showing changes before and after applying different augmentation techniques.

Model	Dataset	Before aug- mentation		Synonym re- placement		Back transla- tion		Contextual augmentation		Adversarial examples		XAI-SR- BT (ours)		XAI-PR- BT (ours)	
		Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
XLMR	Amharic	0.865	0.864	0.872	0.871	<u>0.883</u>	0.882	0.869	0.868	0.870	0.870	0.931	0.931	0.911	0.911
	Arabic	0.775	0.769	0.790	0.788	<u>0.807</u>	0.804	0.778	0.776	0.782	0.781	0.839	0.836	0.814	0.811
	Hindi	0.812	0.810	0.828	0.827	<u>0.843</u>	0.843	0.824	0.823	0.819	0.819	0.887	0.885	0.863	0.860
	Italian	0.689	0.678	0.697	0.693	<u>0.716</u>	0.713	0.692	0.691	0.701	0.700	0.749	0.748	0.733	0.733
	Portuguese	0.712	0.709	0.726	0.724	<u>0.747</u>	0.742	0.718	0.718	0.720	0.719	0.796	0.795	0.766	0.761
	Spanish	0.735	0.735	0.747	0.747	<u>0.756</u>	0.756	0.741	0.740	0.742	0.739	0.809	0.809	0.773	0.771
mBERT	Amharic	0.870	0.872	0.887	0.880	<u>0.894</u>	0.890	0.877	0.870	0.868	0.861	0.925	0.925	0.913	0.912
	Arabic	0.779	0.766	0.792	0.789	<u>0.798</u>	0.797	0.789	0.776	0.782	0.778	0.824	0.823	0.814	0.812
	Hindi	0.845	0.834	0.867	0.859	<u>0.871</u>	0.871	0.855	0.851	0.850	0.847	0.897	0.892	0.881	0.880
	Italian	0.695	0.695	0.708	0.708	<u>0.724</u>	0.719	0.699	0.696	0.700	0.700	0.762	0.762	0.733	0.733
	Portuguese	0.675	0.669	0.686	0.683	<u>0.694</u>	0.694	0.679	0.672	0.680	0.677	0.744	0.742	0.721	0.711
	Spanish	0.765	0.765	0.774	0.774	<u>0.786</u>	0.780	0.768	0.758	0.770	0.761	0.827	0.822	0.803	0.795

Note: XAI-SR-BT refers to XAI Synonym Replacement with Back Translation, and XAI-PR-BT refers to Paraphrasing Replacement with Back Translation.

Table 5

Performance comparison of conventional and XAI-guided context-aware data augmentation techniques on XLM-R and mBERT models across four low-resource languages for the **sentiment analysis task**. The table presents the accuracy (Acc) and F1 score (F1), showing changes before and after applying different augmentation techniques.

Model	Dataset	Before aug- mentation		Synonym re- placement		Back transla- tion		Contextual augmentation		Adversarial examples		XAI-SR- BT (ours)		XAI-PR- BT (ours)	
		Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
XLM-R	Amharic	0.803	0.801	0.834	0.830	0.843	0.842	0.819	0.814	0.828	0.811	0.884	0.883	0.863	0.854
	Hausa	0.827	0.825	0.847	0.844	0.851	0.850	0.837	0.836	0.844	0.844	0.907	0.896	0.875	0.874
	Kinyarwanda	0.706	0.700	0.728	0.724	0.737	0.737	0.710	0.710	0.716	0.715	0.764	0.761	0.749	0.741
	Swahili	0.648	0.626	0.651	0.650	0.660	0.660	0.647	0.643	0.659	0.658	0.687	0.679	0.671	0.663
mBERT	Amharic	0.814	0.811	0.847	0.844	0.855	0.855	0.828	0.822	0.833	0.832	0.892	0.891	0.872	0.870
	Hausa	0.819	0.814	0.854	0.851	0.861	0.860	0.836	0.834	0.840	0.840	0.914	0.913	0.884	0.881
	Kinyarwanda	0.717	0.710	0.731	0.730	0.744	0.745	0.721	0.721	0.724	0.724	0.776	0.774	0.751	0.743
	Swahili	0.661	0.653	0.671	0.670	0.680	0.680	0.665	0.662	0.669	0.668	0.698	0.693	0.689	0.686

Note: XAI-SR-BT refers to XAI Synonym Replacement with Back Translation, and XAI-PR-BT refers to Paraphrasing Replacement with Back Translation.

while introducing controlled variability that enhances the model's generalization ability.

We ensured that critical features remained consistent across datasets, which helped us maintain semantic integrity and prevent distortions that could undermine the model's prediction. When discrepancies in feature importance were detected—such as shifts in the importance of crucial features—we adjusted the augmentation parameter, including the value of k in the $top-k$ least important feature selection parameter, to ensure that the process accurately targeted and preserved relevant features. To ensure meaningful augmentation while preserving semantic integrity, we limited k to 30% of the total features or words in our experiment, selecting them from the least to the most important ranked words based on the input length of our dataset, as shown in Table 1. The value of k was dynamically adjusted for each input text based on its length, with shorter inputs assigned a smaller k and longer inputs assigned a larger k . Additionally, we can generate diverse augmented datasets to optimize performance by varying the value of k and evaluating iteratively. The augmentations are driven by the model's internal understanding of the input. As a result, XAI-guided context-aware data augmentation techniques consistently improve accuracy and F1 scores across several languages, outperforming conventional augmentation methods that often yield limited improvements, as shown in Tables 4 and 5.

The performance differences in XAI-guided data augmentation between low-resource and high-resource languages can be attributed to linguistic characteristics and differences in corpus size. Multilingual models may tend to be biased toward high-resource languages due to factors such as syntactic regularity, lexical richness, and larger corpus sizes. These factors enable high-resource languages to benefit more from augmentation techniques, resulting in greater improvements in model

performance. As shown in Tables 2 and 3, low-resource languages show lower gains in both accuracy and F1 scores. For instance, Swahili and Kinyarwanda achieved only 2.8% and 3.4% accuracy improvements, respectively, on sentiment analysis tasks when using the XAI Paraphrasing with Back Translation (XAI-PR-BT) approach. In contrast, high-resource languages such as Hindi, supported by richer lexical resources and larger corpora, demonstrated significant improvements, for example, a 7.5% gain in accuracy on the hate speech task using the XAI Synonym Replacement with Back Translation (XAI-SR-BT) method.

These performance differences are further influenced by linguistic characteristics such as morphological complexity, syntactic variability, and limited corpus availability. Morphologically rich languages, such as many low-resource Indigenous Australian languages studied in Tosolini and Bowern (2025), often exhibit complex inflectional and derivational systems, which create challenges for data augmentation techniques designed for high-resource languages with simpler morphology. Syntactic diversity, such as the order of free words in some languages, can diminish the effectiveness of text-based augmentation methods that assume more rigid syntactic structures (Li et al., 2022). Additionally, corpus size exacerbates these issues, as low-resource languages lack sufficient data to train robust models, rendering augmentation less effective compared to high-resource languages, where abundant data ensures better generalization (Nzeyimana, 2024; Solyman et al., 2023).

4.1. Comparison of XAI-guided and conventional data augmentation

Tables 4 and 5 compare XAI-guided context-aware data augmentation techniques with conventional methods in multiple languages using the XLM-R and mBERT models. In our study, we compare the

effectiveness of our XAI-guided context-aware data augmentation approach with four conventional techniques, which serve as benchmarks to evaluate our proposed method. Back Translation involves translating the text into another language and then back into the original language to generate variations (Feldman & Coto-Solano, 2020). Random Synonym Replacement with Back Translation substitutes words in the original text with synonyms using translation (Beddiar et al., 2021). Contextual Augmentation employs pre-trained language models, such as mBERT and XLM-R, to replace words based on their context (Kobayashi, 2018). Adversarial examples with back translation introduce small perturbations or noise into the text (Ebrahimi et al., 2017). Across all datasets, as shown in Tables 2 and 3, both XAI-guided context-aware data augmentation methods—XAI Synonym Substitution with Back Translation and XAI Paraphrasing with Back Translation—consistently outperform the baseline results before augmentation and all other conventional augmentation methods. XAI Synonym Substitution with Back Translation provides more significant gains in most cases, showing more substantial improvements in accuracy and F1 score. This highlights the effectiveness of integrating explainability techniques into data augmentation to generate meaningful and diverse text variations that improve model performance across multiple datasets.

Our XAI-guided context-aware data augmentation techniques, such as XAI Synonym Replacement with Back Translation and XAI Paraphrasing with Back Translation, show a clear performance advantage over conventional augmentation methods, including Synonym Replacement, Back Translation, Contextual Augmentation, and Adversarial Examples, as shown in Tables 4 and 5.

Appendix A.3 illustrates examples of augmented data in Spanish and Amharic languages, along with their English translations, by using a synonym replacement with a back-translation approach.

4.2. Evaluation

We utilized two complementary evaluation approaches to determine the effectiveness of our proposed XAI-guided context-aware data augmentation method: model-based and XAI-based evaluation. The model-based approach involves assessing the overall model performance gains before and after augmentation. The XAI-based approach uses XAI techniques to evaluate changes in feature importance and interpretability before and after augmentation.

For the model-based evaluation, we employed XLM-R and mBERT as evaluation models to analyze the performance of both the original and the augmented datasets. By comparing the accuracy and F1-score metrics, we measured the impact of augmented data on model performance. This enabled us to quantify the gains achieved through augmentation regarding generalization and robustness across different language scenarios, as shown in Tables 2 and 3.

The XAI-based evaluation leveraged feature attribution methods to better understand model behavior on the original and augmented datasets. We applied the Integrated gradient XAI technique to determine whether critical features identified in the original data were preserved in the augmented versions. By comparing feature importance scores and XAI explanations before and after augmentation, we assessed the impact of our augmentation strategies on model interpretability.

Our approach is not static but rather an ongoing process of refinement. This iterative feedback loop of model performance evaluation and XAI-based insights allows us to refine the augmentation process continuously. By doing so, we optimize model accuracy and robustness and enhance interpretability, ensuring that our augmented data effectively supports the model prediction performance and explainability goals.

5. Limitations and future direction

Despite the strong performance gains demonstrated by our XAI-guided context-aware data augmentation approach, some limitations

remain. The method depends on the selected explainability technique. This may introduce computational overhead due to the integration of XAI methods into the baseline model. Moreover, the reliability of current XAI explanations can be inconsistent, and the quality of augmentation may vary depending on the chosen XAI method and translation tool. In future work, we plan to investigate more efficient and stable XAI techniques and extend our evaluation across a broader range of model architectures (language-specific models like AfriBERTa and AraBERT), languages, and tasks to analyze performance improvements. We also intend to explore the implications of our approach for multilingual fairness and adversarial robustness.

6. Conclusion

Our study introduces a novel XAI-guided context-aware data augmentation approach that effectively addresses the limitations of conventional data augmentation methods, particularly by improving the performance of language models in different languages. The proposed approach ensures that critical semantic information is preserved during augmentation by integrating XAI techniques with multilingual models to identify less model-influential features in the input text. This targeted strategy results in high-quality augmented data, significantly enhancing various language models' robustness, performance, and explainability, offering a promising future, particularly for low-resource language processing. We demonstrate that XAI-guided context-aware data augmentation enhances the generalization of language models and provides transparency into the augmentation process. Through iterative refinement using feedback loops between XAI and applied models, the approach maintains contextual integrity, making it particularly suitable for addressing the diverse challenges of several languages. Our experimental results confirm that the XAI-guided context-aware data augmentation approach consistently outperforms existing or conventional augmentation methods, achieving higher accuracy and F1 scores across multiple language datasets. By integrating explainability techniques and language models into the data augmentation, we bridge the gap between enhancing model performance and maintaining interpretability. Our study sets the stage for the broader application of XAI in enhancing various language models in different languages.

CRedit authorship contribution statement

Melkamu Abay Mersha: Conceptualization, Methodology, Investigation, Formal analysis, Writing - review & editing; **Mesay Gemedo Yigezu:** Validation, Data curation, Visualization; **Atnafu Lambebo Tonja:** Writing - review & editing; **Hassan Shakil:** Methodology, Formal analysis, Writing - review & editing; **Samer Iskander:** Writing - review & editing, Visualization; **Olga Kolesnikova:** Resources, Supervision; **Jugal Kalita:** Supervision, Project administration, Writing - review & editing.

Data availability

The authors do not have permission to share data.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Melkamu Abay Mersha reports was provided by University of Colorado Colorado Springs. Melkamu Abay Mersha reports a relationship with University of Colorado Colorado Springs that includes: non-financial support. Melkamu Abay Mersha has patent pending to No. No. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Sample augmented data A.3

Lang.	Original Text	Augmented Text	Original Translation	Augmented Translation
Amharic	እነዚህ ወልጋዳ ሰዎች በ ሽማግሌ በተመሰለው እንቁ ሰው ወንድማዊ ምክር ወደ መልካሙ ነገር ሊመጡ መንገዳቸውን ሊፈትሹ ሊያስተክክሉ ይገባል	እነዚህ ወልጋዳ ሰዎች በ ሽማግሌ በተመሰለው ክቡር ሰው ወንድማዊ ምክር ወደ ጥሩ ነገር ሊመጡ መንገዳቸውን ሊመረምሩ ሊያስተክክሉ ይገባል	These Wolgada (bad) people should check their ways and come to the good thing with the brotherly advice of the gem person who is like an old man.	These Wolgada (bad) people should inspect their ways to come to the great thing with the brotherly advice of a noble person who is like an old man.
	ስከፍት በመጀመሪያ ማንበብ የምፈልገው ምርጥ ፀሀፊ ምርጥ ተንታኝ ምርጥ ኢኮኖሚስት ምርጥ አንባቢ ብቻ እረጅም እድሜ እመኝልሁለሁ ብዘገይም መልካም ልደት	ስጀምር በቅድሚያ ማንበብ የምፈልገው ምርጥ ፀሀፊ ምርጥ ተንታኝ ምርጥ ኢኮኖሚስት ምርጥ አንባቢ ብቻ እረጅም እድሜ እመኝልሁለሁ ብዘገይም መልካም ልደት	When I open it, the first thing I want to read is the best writer, the best analyst, the best economist, the best reader. I wish you a long life, even if I'm late.	When I start it, the primary thing I want to read is the best writer, the best analyst, the best economist, the best reader. I wish you a long life, even if I'm late.
	ወድሜ አምላክ ዘመንህ ህሊናህ ስራህ ይባርክ በዚህ ጉዳይ ከብዙ ሰዎች ጋር አወጋው ነበር ለከ ቀድመው የሂዱ ወንድሞች አሉኝ አግኝቼ ባወራህ እንዴት ደሰ ባለኝ	ወድሜ አምላክ ዕድሜህን ህሊናህ ተግባርህን ይባርክ በዚህ ጉዳይ ከብዙ ሰዎች ጋር አወራው ነበር ለከ ቀድመው የሂዱ ወንድሞች አሉኝ አግኝቼ ባወራህ እንዴት ደሰ ባለኝ	Brother, God bless your life, your conscience, your work. I have discussed this matter with many people. I have brothers who have gone before me.	Brother, God bless your age, your conscience, your job. I have talked this matter with many people. I have brothers who have gone before me.
Spanish	Trump urge al Congreso aprobar recursos para construir el muro "más rápido". Pues que contraten mexicanos indocumentados para levantarlo... y al final que los deporten.	Trump urge al Congreso aprobar fondos para construir el muro "más rápido". Pues que contraten mexicanos indocumentados para levantarlo... además al término que los expulsen	Trump urges Congress to approve resources to build the wall "faster." Well, let them hire undocumented Mexicans to raise it... and in the end, let them deport them.	Trump urges Congress to approve funds to build the wall "faster." Well, let them hire undocumented Mexicans to build it... and at the end they will expel them
	Es flipante lo que ha dicho, Es la típica persona que justifica cualquier violación o acoso. Pero lo más triste es que puede que detrás de su pj esté un due mujer.	Es alucinante lo que ha dicho, Es la típica persona que justifica cualquier violación o acoso. Sin embargo, lo más lamentable es que podría que detrás de su pj esté un supuesto mujer.	What he said is amazing. He is the typical person who justifies any violation or harassment. But the saddest thing is that there may be a female due behind his pj.	What he said is amazing. He is the typical person who justifies any violation or harassment. However, the unpleasant thing is that there could be a woman behind his pj.
	A usted le debo un sueño que fue ilusión, a usted le debo muchos besos por las noches y mi lengua debe un recorrido en la ruta de sus lunares... a usted le debo mucho, y no pienso quedar en deuda.	A usted le debo un sueño que fue fantasía, a usted le debo varios besos por las noches y mi lengua debe un trayecto en el camino de sus lunares... a usted le debo demasiado, y no pienso permanecer en deuda.	I owe you a dream that was an illusion, I owe you many kisses at night and my tongue owes a journey along the route of your moles... I owe you a lot, and I do not plan to remain in debt.	I owe you a dream that was fantasy, I owe you several kisses at night and my tongue owes a journey along the path of your moles... I owe you too much, and I do not intend to remain in debt.

Fig. A.3. Sample Spanish and Amharic augmented data with translations from the Hate Speech dataset using the Synonym Replacement with Back Translation Approach. These samples include representative hate speech examples, which may contain offensive language, solely for research and analysis in the context of hate speech detection and data augmentation.

References

- Abnar, S., & Zuidema, W. (2020). Quantifying attention flow in transformers. In *Proceedings of the 58th annual meeting of the association for computational linguistics* (pp. 4190–4197).
- Alzantot, M., Sharma, Y., Elgohary, A., Ho, B.-J., Srivastava, M., & Chang, K.-W. (2018). Generating natural language adversarial examples. *arXiv preprint arXiv:1804.07998*.
- Ansari, G., Kaur, P., & Saxena, C. (2024). Data augmentation for improving explainability of hate speech detection. *Arabian Journal for Science and Engineering*, 49(3), 3609–3621.
- Atanasova, P. (2024). A diagnostic study of explainability techniques for text classification. In *Accountable and explainable methods for complex reasoning over text* (pp. 155–187). Springer.
- Ayele, A. A., Yimam, S. M., Belay, T. D., Asfaw, T., & Biemann, C. (2023). Exploring amharic hate speech data collection and classification approaches. In *Proceedings of the 14th international conference on recent advances in natural language processing* (pp. 49–59).
- Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.-R., & Samek, W. (2015). On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS One*, 10(7), e0130140.
- Basile, V., Bosco, C., Fersini, E., Nozza, D., Patti, V., Pardo, F. M. R., Rosso, P., & Sanguinetti, M. (2019). Semeval-2019 task 5: Multilingual detection of hate speech against immigrants and women in twitter. In *Proceedings of the 13th international workshop on semantic evaluation* (pp. 54–63).
- Bayer, M., Kaufhold, M.-A., Buchhold, B., Keller, M., Dallmeyer, J., & Reuter, C. (2023). Data augmentation in natural language processing: A novel text generation approach for long and short text classifiers. *International Journal of Machine Learning and Cybernetics*, 14(1), 135–150.
- Bayer, M., Kaufhold, M.-A., & Reuter, C. (2022). A survey on data augmentation for text classification. *ACM Computing Surveys*, 55(7), 1–39.
- Bhattacharya, A. Applied machine learning explainability techniques: make ML models explainable and trustworthy for practical applications using LIME, SHAP, and more, (2022). Packt Publishing Ltd..
- Beddier, D. R., Jahan, M. S., & Oussalah, M. (2021). Data expansion using back translation and paraphrasing for hate speech detection. *Online Social Networks and Media*, 24, 100153.
- Bhardwaj, M., Akhtar, M. S., Ekbal, A., Das, A., & Chakraborty, T. (2020). Hostility detection dataset in hindi. *arXiv preprint arXiv:2011.03588*.
- Bird, S., Klein, E., Loper, S., Natural language processing with Python: analyzing text with the natural language toolkit, O'Reilly Media, Inc., (2019).
- Bosco, C., Dell'Orletta, F., Poletto, F., Sanguinetti, M., Tesconi, M. et al. (2018). Overview of the evalita 2018 hate speech detection task. In *Ceur workshop proceedings* (pp. 1–9). CEUR (vol. 2263).
- Chefer, H., Gur, S., & Wolf, L. (2021). Transformer interpretability beyond attention visualization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 782–791).
- Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., & Stoyanov, V. (2019). Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.
- Corbett, J.-P., & Ghadivel, H. A. (2020). Bet: A backtranslation approach for easy data augmentation in transformer-based paraphrase identification context. *arXiv preprint arXiv:2009.12452*.
- Dai, H., Liu, Z., Liao, W., Huang, X., Cao, Y., Wu, Z., Zhao, L., Xu, S., Zeng, F., Liu, W. et al. (2025). AugGPT: Leveraging ChatGPT for Text Data Augmentation. *IEEE Transactions on Big Data* 11(3), 665–678.
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Ding, B., Qin, C., Zhao, R., Luo, T., Li, X., Chen, G., Xia, W., Hu, J., Tuan, L. A., & Joty, S. (2024). Data augmentation using llms: Data perspectives, learning paradigms and challenges. In *Findings of the association for computational linguistics ACL 2024* (pp. 1679–1705).
- Ebrahimi, J., Rao, A., Lowd, D., & Dou, D. (2017). Hotflip: White-box adversarial examples for text classification. *arXiv preprint arXiv:1712.06751*.
- Fantozzi, P., & Naldi, M. (2024). The explainability of transformers: Current status and directions. *Computers*, 13(4), 92.
- Feldman, I., & Coto-Solano, R. (2020). Neural machine translation models with back-translation for the extremely low-resource indigenous language bribri. In *Proceedings of the 28th international conference on computational linguistics* (pp. 3965–3976).
- Fellbaum, C. (1998). WordNet: An electronic lexical database. MIT Press.
- Fortuna, P., da Silva, J. R., Wanner, L., Nunes, S. et al. (2019). A hierarchically-labeled portuguese hate speech dataset. In *Proceedings of the third workshop on abusive language online* (pp. 94–104).
- Hadi, M. U., Al-Tashi, Q., Qureshi, R., Shah, A., Muneer, A., Irfan, M., Zafar, A., Shaikh, M. B., Akhtar, N., Al-Garadi, M. A. et al. Lms: A comprehensive survey of applications, challenges, datasets, models, limitations, and future prospects.
- Hayashi, T., Watanabe, S., Zhang, Y., Toda, T., Hori, T., Astudillo, R., & Takeda, K. (2018). Back-translation-style data augmentation for end-to-end ASR. In *2018 IEEE spoken language technology workshop (SLT)* (pp. 426–433). IEEE.
- Holzinger, A., Saranti, A., Molnar, C., Biecek, P., & Samek, W. (2022). Explainable AI methods-a brief overview. In *International workshop on extending explainable AI beyond deep models and classifiers* (pp. 13–38). Springer.
- Ismail, A. A., Corrada Bravo, H., & Feizi, S. (2021). Improving deep learning interpretability by saliency guided training. *Advances in Neural Information Processing Systems*, 34, 26726–26739.
- Iyyer, M., Manjunatha, V., Boyd-Graber, J., & Daumé III, H., (2015). Deep unordered composition rivals syntactic methods for text classification. In *Proceedings of the 53rd annual meeting of the association for computational linguistics and the 7th international joint conference on natural language processing (volume 1: Long papers)* (pp. 1681–1691).
- Kapishnikov, A., Venugopalan, S., Avci, B., Wedin, B., Terry, M., & Bolukbasi, T. (2021). Guided integrated gradients: An adaptive path method for removing noise. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 5050–5058).
- Kapusta, J., Držić, D., Šteflovíč, K., & Nagy, K. S. (2024). Text data augmentation techniques for word embeddings in fake news classification. *IEEE Access*, 12, 31538–31550.
- Khan, R., & Venugopal, A. (2024). Exploring deep learning methods for text augmentation to handle imbalanced datasets in natural language processing. In *2024 3rd edition of IEEE delhi section flagship conference (DELCON)* (pp. 1–8). IEEE.
- Kim, B., Wattenberg, M., Gilmer, J., Cai, C., Wexler, J., Viegas, F. et al. (2018). Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). In *International conference on machine learning* (pp. 2668–2677). PMLR.
- Kobayashi, S. (2018). Contextual augmentation: Data augmentation by words with paradigmatic relations. *arXiv preprint arXiv:1805.06201*.
- Kumar, V., Choudhary, A., & Cho, E. (2020). Data augmentation using pre-trained transformer models. *arXiv preprint arXiv:2003.02245*.
- Kwon, S., & Lee, Y. (2023). Explainability-based mix-up approach for text data augmentation. *ACM Transactions on Knowledge Discovery from Data*, 17(1), 1–14.
- Latief, A. D., Jarin, A., Yaniasih, Y., Afrā, D. I. N., Nurfadhilah, E., Pebiana, S., Hidayati, N. N., & Fajri, R. (2024). Latest research in data augmentation for low resource language text translation: A review. In *2024 international conference on computer, control, informatics and its applications (IC3INA)* (pp. 185–190). IEEE.
- Li, B., Hou, Y., & Che, W. (2022). Data augmentation approaches in natural language processing: A survey. *Ai Open*, 3, 71–90.
- Liu, Y. (2019). Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Lundberg, S. (2017). A unified approach to interpreting model predictions. *arXiv preprint arXiv:1705.07874*.
- Mersha, M., Bitewa, M., Abay, T., & Kalita, J. (2024a). Explainability in neural networks for natural language processing tasks. *arXiv preprint arXiv:2412.18036*.
- Mersha, M., Lam, K., Wood, J., AlShami, A., & Kalita, J. (2024b). Explainable artificial intelligence: A survey of needs, techniques, applications, and future direction. *Neuro-computing* 599, 128111.
- Mersha, M. A., Kalita, J. et al. (2024c). Semantic-driven topic modeling using transformer-based embeddings and clustering algorithms. *Procedia Computer Science*, 244, 121–132.
- Mersha, M. A., Yigezu, M. G., & Kalita, J. (2025). Evaluating the effectiveness of XAI techniques for encoder-based language models. *Knowledge-Based Systems* 310 (p.113042).
- Muhammad, S., Abdulmumin, I., Ayele, A., Ousidhoum, N., Adelani, D., Yimam, S., Ahmad, I., Beloucif, M., Mohammad, S., Ruder, S., Hourrane, O., Jorge, A., Brazdil, P., Ali, F., David, D., Osei, S., Shehu-Bello, B., Lawan, F., Gwadabe, T., ...Belay, T., (2023). AfriSenti: A twitter sentiment analysis benchmark for African languages. In H. Bouamor, J. Pino, & K. Bali (Eds.), *Proceedings of the 2023 conference on empirical methods in natural language processing* (pp. 13968–13981). Singapore: Association for Computational Linguistics. <https://aclanthology.org/2023.emnlp-main.862>. <https://doi.org/10.18653/v1/2023.emnlp-main.862>
- Nair, A. R., Singh, R. P., Gupta, D., & Kumar, P. (2024). Evaluating the impact of text data augmentation on text classification tasks using distilBERT. *Procedia Computer Science*, 235, 102–111.
- Nishi, K., Ding, Y., Rich, A., & Hollerer, T. (2021). Augmentation strategies for learning with noisy labels. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 8022–8031).
- Nzeyimana, A. (2024). Low-resource neural machine translation with morphological modeling. *arXiv preprint arXiv:2404.02392*.
- Ousidhoum, N., Lin, Z., Zhang, H., Song, Y., & Yeung, D.-Y. (2019). Multilingual and multi-aspect hate speech analysis. *arXiv preprint arXiv:1908.11049*.
- Pires, T., Schlinger, E., & Garrette, D. (2019). How multilingual is multilingual BERT? *arXiv preprint arXiv:1906.01502*.
- Qiang, Y., Pan, D., Li, C., Li, X., Jang, R., & Zhu, D. (2022). Atcat: Explaining transformers via attentive class activation tokens. *Advances in Neural Information Processing Systems*, 35, 5052–5064.
- Radford, A., Narasimhan, K., Salimans, T., Sutskever, I. et al. (2018). Improving language understanding by generative pre-training.
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should i trust you?” Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 1135–1144).
- Schick, T., & Schütze, H. (2021). Generating datasets with pretrained language models. *arXiv preprint arXiv:2104.07540*.
- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2020). Grad-CAM: Visual explanations from deep networks via gradient-based localization. *International Journal of Computer Vision*, 128, 336–359.
- Sennrich, R. (2015). Improving neural machine translation models with monolingual data. *arXiv preprint arXiv:1511.06709*.
- Serrano, S., & Smith, N. A. (2019). Is attention interpretable? *arXiv preprint arXiv:1906.03731*.
- Shakil, H., Mahi, A. M., Nguyen, P., Ortiz, Z., & Mardini, M. T. (2024). Evaluating text summaries generated by large language models using openAI's GPT. *arXiv preprint arXiv:2405.04053*.
- Shorten, C., Khoshgoftaar, T. M., & Furht, B. (2021). Text data augmentation for deep learning. *Journal of Big Data*, 8(1), 101.
- Shrikumar, A., Greenside, P., & Kundaje, A. (2017). Learning important features through propagating activation differences. In *International conference on machine learning* (pp. 3145–3153). PMLR.
- Solyman, A., Zappatore, M., Zhenyu, W., Mahmoud, Z., Alfatemi, A., Ibrahim, A. O., & Gabralla, L. A. (2023). Optimizing the impact of data augmentation for low-resource

- grammatical error correction. *Journal of King Saud University-Computer and Information Sciences*, 35(6), 101572.
- Srinivas, S., & Fleuret, F. (2019). Full-gradient representation for neural network visualization. *Advances in Neural Information Processing Systems*, 32. 33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada
- Sundararajan, M., Taly, A., & Yan, Q. (2017). Axiomatic attribution for deep networks. In *International conference on machine learning* (pp. 3319–3328). PMLR.
- Taheri, A., Zamanifar, A., & Farhadi, A. (2024). Enhancing aspect-based sentiment analysis using data augmentation based on back-translation. *International Journal of Data Science and Analytics* 19, 1–26.
- Takahagi, K., & Shinnou, H. (2023). Data augmentation by shuffling phrases in recognizing textual entailment. In *Proceedings of the 37th Pacific Asia conference on language, information and computation* (pp. 194–200).
- Tonja, A. L., Mersha, M., Kalita, A., Kolesnikova, O., & Kalita, J. (2023). First attempt at building parallel corpora for machine translation of northeast india's very low-resource languages. *arXiv preprint arXiv:2312.04764*
- Tosolini, A., & Bowern, C. (2025). Data augmentation and hyperparameter tuning for low-resource MFA. *arXiv preprint arXiv:2504.07024*
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., Polosukhin, I. (2017). Attention is All You Need. *Advances in Neural Information Processing Systems* 30.
- Vig, J. (2019). A multiscale visualization of attention in the transformer model. *arXiv preprint arXiv:1906.05714*
- Volpi, R., Namkoong, H., Sener, O., Duchi, J. C., Murino, V., & Savarese, S. (2018). Generalizing to unseen domains via adversarial data augmentation. *Advances in Neural Information Processing Systems*, 31. 32nd Conference on Neural Information Processing Systems (NeurIPS 2018), Montréal, Canada.
- Wang, K., Zhu, J., Ren, M., Liu, Z., Li, S., Zhang, Z., Zhang, C., Wu, X., Zhan, Q., Liu, Q. et al. (2024). A survey on data synthesis and augmentation for large language models. *arXiv preprint arXiv:2410.12896*
- Wei, J., & Zou, K. (2019). Eda: Easy data augmentation techniques for boosting performance on text classification tasks. *arXiv preprint arXiv:1901.11196*
- Yigezu, M. G., Mersha, M. A., Bade, G. Y., Kalita, J., Kolesnikova, O., & Gelbukh, A. (2024). Ethio-fake: Cutting-edge approaches to combat fake news in under-resourced languages using explainable ai. *arXiv preprint arXiv:2410.02609*
- Yuan, T., Li, X., Xiong, H., Cao, H., & Dou, D. (2021). Explaining information flow inside vision transformers using markov chain. In *Explainable AI approaches for debugging and diagnosis*.
- Zeiler, M. D. (2014). Visualizing and understanding convolutional networks. In *European conference on computer vision/arxiv*. (vol. 1311).
- Zhang, J., Zhao, Y., Saleh, M., & Liu, P. (2020). Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International conference on machine learning* (pp. 11328–11339). PMLR.
- Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., & Torralba, A. (2016). Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 2921–2929).
- Zhuo, Y., Ge, Z. (2024). IG2: Integrated Gradient on Iterative Gradient Path for Feature Attribution. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 46(11).