

Enhancing an RNN-Attention Yoruba Text Autocompletion System through an Optimized Adam Framework

Oluokun Samuel Olugbenga^{1*}, Ayeni Joshua Ayobami², Adebunmi Adefunso³, Oladayo Ezekiel Makinde⁴, Adebayo Ademola Riliwan⁵

¹Department of Information Systems and Technology, Kings University, Odeomu, Nigeria

²Department of Computer Science, Ajayi Crowther University, Oyo, Nigeria

³Athenahealth, Boston Massachusetts, USA

⁴Department of Computer Science, Ajayi Crowther University, Oyo, Nigeria

⁵Federal School of Surveying, Oyo, Nigeria

DOI: <https://doi.org/10.51584/IJRIAS.2025.10100000170>

Received: 21 October 2025; Accepted: 28 October 2025; Published: 20 November 2025

ABSTRACT

The development of effective neural models for low-resource languages is fundamentally constrained by two interrelated factors: architectural suitability for linguistic complexity and optimization stability on small datasets. This research addresses the critical yet under-explored challenge of optimization instability for character-level sequence modeling in Yoruba, a morphologically rich and tonal language. We posit that standard adaptive optimizers like Adam, while performant in high-resource contexts, introduce convergence pathologies in low-resource settings due to volatile gradient estimates and an inability to adapt to sparse loss landscapes. To address this, we propose a principled enhancement to the Adam optimizer, integrating a dynamic learning rate scheduler, gradient norm clipping, and a strategically determined batch size. This Enhanced Adam framework is applied to a character-level Recurrent Neural Network augmented with a multi-head attention mechanism, an architecture designed to handle Yoruba's agglutinative and tonal features. In a rigorous comparative study, the model trained with our Enhanced Adam optimizer achieved a perplexity of 2.07, a statistically significant 8.5% improvement over the identical architecture trained with standard Adam (perplexity 2.26). More importantly, the enhanced framework demonstrably improved training stability, accelerated convergence, and yielded a better-calibrated model. This work establishes that targeted optimizer engineering is not merely an implementation detail but a critical research direction for unlocking the full potential of advanced neural architectures in low-resource Natural Language Processing (NLP), providing a reproducible and transferable methodology for other underserved languages.

Keywords: Low-Resource NLP, Yoruba Language, Text Autocompletion, Adam Optimizer, Optimization Stability, Gradient Clipping, Learning Rate Scheduling, RNN, Attention Mechanism.

INTRODUCTION

The transformative advances in Natural Language Processing (NLP) over the past decade, driven by deep learning, have predominantly served a handful of high-resource languages such as English, Mandarin, and Spanish (Joshi *et al.*, 2020). This has created a significant digital divide, leaving speakers of thousands of other languages without access to foundational technologies like accurate machine translation, robust speech recognition, and intelligent writing assistants (Blasi *et al.*, 2022). Among these underserved languages is Yoruba, a major Niger-Congo language spoken by over 40 million people in West Africa and the diaspora. The lack of

NLP tools for Yoruba impedes digital inclusion, hinders educational and economic opportunities, and contributes to the erosion of linguistic diversity in the digital sphere (Adelani *et al.*, 2021).

Developing NLP tools for a language like Yoruba presents a dual challenge. The first challenge is architectural: designing models that can effectively capture the language's unique linguistic characteristics. Yoruba is tonal, where meaning is lexically determined by pitch patterns on vowels (e.g., *igbá* calabash vs. *igbà* time), and agglutinative, forming complex words through the linear combination of morphemes (Adeniyi, 2020; Akinola *et al.*, 2021). These properties make sub-word or character-level modeling more appropriate than word-level modeling, as the latter fails to capture the nuanced, compositional structure of words (Bostrom & Durrett, 2020).

The second, often understated challenge, is optimization. State-of-the-art neural architectures, often comprising millions of parameters, are data-hungry. In low-resource settings, the scarcity of data leads to a poorly defined optimization landscape. Standard optimization algorithms, most notably the Adam optimizer (Kingma & Ba, 2015), which is ubiquitous in modern deep learning, were designed and tuned for large-scale datasets. When applied to small corpora, Adam and its variants can exhibit pathological behaviors: unstable convergence due to noisy gradient estimates, vulnerability to exploding gradients in deep recurrent networks, and an inability to escape sharp local minima due to a fixed or poorly adapted learning rate schedule (Zhang *et al.*, 2020; Wilson *et al.*, 2017). Consequently, a powerful architecture may never reach its potential because the training process is inherently unstable and inefficient.

While previous work, including our own, has successfully demonstrated the efficacy of attention-augmented RNNs for Yoruba text autocompletion (Oluokun *et al.*, 2025), the optimization process itself was treated as a static component. This paper directly addresses this gap by making optimization the central subject of inquiry. We hypothesize that a purpose-built optimization framework is essential to stabilize training and maximize the performance of complex models on low-resource data.

The primary contributions of this research are:

1. A detailed analysis of the optimization challenges specific to training character-level RNN-Attention models on a small Yoruba corpus.
2. The design and implementation of an Enhanced Adam optimization framework that systematically incorporates dynamic learning rate scheduling, gradient clipping, and strategic batch processing to mitigate these challenges.
3. An empirical evaluation demonstrating that the proposed framework yields a statistically significant improvement in model performance (perplexity) and training stability compared to the standard Adam optimizer, using an identical neural architecture.
4. The provision of a robust, reproducible methodology that can be adapted for developing NLP systems for other low-resource languages.

This work argues that for the field of low-resource NLP to mature, optimizer engineering must be recognized as a discipline as critical as architectural innovation.

LITERATURE REVIEW

2.1 NLP for Low-Resource and Morphologically Rich Languages

The plight of low-resource languages in NLP has garnered increasing attention. Initiatives like MasakhaNEWS (Adelani *et al.*, 2023) for news classification and MasakhaNER (Adelani *et al.*, 2021) for named entity recognition have created crucial benchmarks for African languages. For Yoruba specifically, research has focused on diacritic restoration (Ogheneruemu *et al.*, 2023; Akindele *et al.*, 2024), part-of-speech tagging (Ugwu *et al.*, 2024), and machine translation (Dossou & Emezue, 2021). A common thread is the need for models that operate

at the sub-word level. Bostrom & Durrett (2020) showed that character-level models often outperform word-level models on morphologically complex tasks because they can handle out-of-vocabulary words and model morphological processes directly. This justifies our choice of a character-level modeling approach for Yoruba autocompletion.

2.2 Neural Architectures for Sequence Prediction

Recurrent Neural Networks (RNNs), and specifically Long Short-Term Memory (LSTM) networks (Hochreiter & Schmidhuber, 1997), have been the historical standard for sequence prediction tasks like text generation and autocompletion (Sutskever *et al.*, 2011). However, their sequential nature and tendency to "forget" information from the distant past due to vanishing gradients limit their effectiveness. The attention mechanism (Bahdanau *et al.*, 2015; Vaswani *et al.*, 2017) revolutionized sequence modeling by allowing the model to directly reference any part of the input sequence, effectively creating a dynamic, content-addressable memory. While Transformers have largely superseded RNNs in high-resource settings, their data efficiency is poor. Therefore, an RNN augmented with attention presents a compelling middle ground: it retains the inductive bias for sequential processing while gaining the representational power of attention, making it potentially more suitable for low-resource scenarios (Al-Anzi & Shalini, 2024; Vanama *et al.*, 2023).

2.3 Optimization in Deep Learning

The Adam optimizer (Kingma & Ba, 2015) is an adaptive learning rate algorithm that combines the advantages of AdaGrad (Duchi *et al.*, 2011) and RMSProp (Tieleman & Hinton, 2012). It maintains per-parameter learning rates based on estimates of the first and second moments of the gradients. While its adaptive nature leads to rapid initial convergence, it has known limitations. Wilson *et al.* (2017) argued that adaptive methods often generalize worse than stochastic gradient descent (SGD) with momentum because they can converge to sharp minima. This issue is exacerbated with small datasets where the gradient noise is high.

While Adam's adaptive nature facilitates rapid initial convergence, its performance, particularly in low-resource scenarios, is highly sensitive to its configuration and can be prone to instability and poor generalization (Wilson *et al.*, 2017). To mitigate these issues, two key families of techniques are essential for stabilizing Adam in practice:

Gradient Clipping: This technique is critical for preventing the exploding gradient problem, which is especially prevalent in recurrent architectures (Pascanu *et al.*, 2013). By constraining the L2-norm of the gradient vector to a predefined threshold (e.g., $\text{clip_value} = 1.0$), it ensures that parameter updates remain within a stable range, preventing destructive large steps that can derail the optimization process. The operation is defined as $\|g\|_2 = \min(\|g\|_2, \text{clip_value})$.

Learning Rate Scheduling and Regularization: A fixed learning rate can hinder convergence to a precise minimum. Dynamic scheduling, such as the Reduce LR On Plateau scheduler which reduces the learning rate by a factor (e.g., 0.5) upon a loss plateau, refines the optimization trajectory in later stages. Furthermore, integrating weight decay as an L2 regularization term directly into the update rule ($\lambda \theta_t \lambda \theta_t$) is a principled method to prevent overfitting and encourage simpler models by penalizing large weights, leading to better generalization.

However, most research on optimizer behavior is conducted in high-resource contexts. The specific failure modes of Adam and the precise calibration of stabilization techniques for low-resource, character-level NLP remain an open area for investigation, which this paper directly addresses.

2.4 Related work

Akindele *et al.*, (2024) study aimed to address the lack of a standard benchmark for evaluating Yoruba diacritization systems and improve automatic diacritization using lightweight models. Yoruba, a tonal language, relies heavily on diacritics for meaning and pronunciation, making accurate diacritization essential. Manual diacritization is time-consuming, necessitating automated solutions. The researchers introduced the Yoruba

Automatic Diacritization (YAD) dataset, derived from MENYO-20k, and pre-trained T5 models (Oyo-T5) specifically for Yoruba. They compared these models with multilingual T5 variants (MT5, AfriMT5, AfriTeVaV2, UMT5) to evaluate performance. The methodology involved pre-training Oyo-T5 models of varying sizes (tiny to base) and fine-tuning them on YAD, Bible, and JW300 datasets. The models were evaluated using SacreBLEU and ChrF metrics on YAD, Global Voices (GV), and Bible test sets. Results showed that Oyo-T5base outperformed larger multilingual models, and increasing model size and training data improved performance. Notably, Oyo-T5-small (60M parameters) surpassed AfriTeVa-base (313M parameters). The study concluded that more data and larger models enhance Yoruba diacritization, with Oyo-T5-base achieving the best results. The YAD dataset and models were released on GitHub for reproducibility.

Ahia *et al.*, (2024) study addressed the limitations in Yoruba natural language processing (NLP) by developing resources and models for regional Yoruba dialects. Despite Yoruba having over 47 million speakers, existing natural language processing research focused primarily on the standard dialect, neglecting regional variations. This gap led to disparities in automatic speech recognition (ASR), machine translation (MT), and speech-to-text translation (S2TT) performance for non-standard dialects. The research sought to create a high-quality parallel corpus, YORULECT, covering four Yoruba dialects, Standard Yoruba, Ifè, Ilàje, and Ìjèbù across religious, news, and TED Talk domains. It aimed to evaluate the zero-shot performance of state-of-the-art natural language processing models on these dialects and fine-tune them to improve performance. Native speakers were engaged to curate a dataset comprising 1,506 parallel text sentences per dialect and 9 hours of recorded speech. Preprocessing involved text normalization, phonetic transcription, and segmentation. Speech data were recorded in sound-isolated environments. Standard natural language processing models like NLLB-600M (for MT), MMS and Whisper (for ASR), and SeamlessM4T (for S2TT) were tested, and dialect-specific fine-tuning was applied. Performance evaluation revealed that standard Yoruba outperformed regional dialects across all tasks, highlighting a lack of robustness in existing models. Zero-shot MT results showed Google Translate performed best but had a significant performance gap between standard and regional Yoruba. After fine-tuning, BLEU scores improved by 14 points, and ASR word error rates decreased by 20 points. S2TT remained the most challenging, with limited improvement post-finetuning.

METHODOLOGY

3.1 Experimental Design and Research Questions

This research employs a controlled, comparative experimental design. The independent variable is the optimization strategy, and the dependent variables are the model's final performance (perplexity, accuracy) and training dynamics (loss convergence, stability). The neural architecture a character-level RNN with a multi-head attention mechanism is held constant across experiments to isolate the effect of the optimizer.

The research seeks to answer the following questions:

1. RQ1: What are the characteristic signs of optimization instability when training an RNN-Attention model on a low-resource Yoruba dataset with the standard Adam optimizer?
2. RQ2: To what extent does the proposed Enhanced Adam framework mitigate these instability issues, as measured by training loss curves and variance?
3. RQ3: Does the Enhanced Adam framework lead to a statistically significant improvement in the final model's predictive performance and confidence on a held-out test set?

3.2. Data Collection and Preprocessing

Due to the lack of a large-scale digital corpus for Yoruba, a custom dataset was curated for this study. Both models were trained on the same dataset extracted from "fdata.xlsx", containing 4,431 words with their accented and non-accented versions. The dataset was downloaded <https://www.kaggle.com/datasets/adeyemiquadri/new-yoruba-data>. The dataset was split into training (80%), validation (10%), and test (10%) sets. The vocabulary consists of 22 unique characters, and the maximum sequence length is 11 characters.

Figure 3.1: Research Dataset

	A	B
1	words with accent	words without accents
2	ibérépépé	iberepepe
3	ihin	ihin
4	rere	rere
5	nipa	nipa
6	jésù	jesu
7	kristi	kristi
8	ganan	ganan
9	gégé	gege
10	bi	bi
11	a	a
12	ti	ti
13	kòwé	kowe
14	rà	ra
15	ninú	ninu
16	àisàyà	aisaya
17	wòlì	wolii
18	pé	pe
19	wó	wo
20	ó	o
21	èmi	emi
22	yóò	yoo
23	rán	ran
24	onṣẹ	onse
25	mi	mi
26	jáde	jade

3.3 Data Preprocessing

The preprocessing pipeline (Figure 3.1) was designed to preserve Yorùbá's phonological and orthographic integrity while converting text into a numerical representation. It consisted of the following steps:"

3.3.1. Text Normalization and Cleaning

The raw text corpus underwent a rigorous normalization process to ensure consistency and eliminate noise. This involved:

Diacritic Preservation: All Yorùbá-specific diacritical marks (e.g., ẹ, ọ, ẹ, à, è, ì, ò, ù) were meticulously preserved, as they are phonemically critical and determine lexical meaning.

Noise Removal: Non-linguistic artifacts, including numerical digits, punctuation marks (except for relevant sentence delimiters used in sequence creation), and extraneous whitespace characters were systematically removed.


Case Normalization: All text was converted to lowercase to maintain a consistent vocabulary and reduce sparsity, a standard practice in character-level modeling.

3.3.2 Character-Level Tokenization

Given Yorùbá's agglutinative morphology, where words are formed by combining morphemes, character-level tokenization was explicitly chosen over word-level tokenization. This approach allows the model to learn

subword morphological units and generate novel, valid words not present in the training data. The tokenization process segmented the normalized text into its constituent characters, treating each character, including spaces and diacritical marks, as a discrete token. For example, the phrase "È kú ààrò" was decomposed into the sequence: ['È', ' ', 'k', 'ú', ' ', 'à', 'à', 'r', 'ò', ' ']. A vocabulary of 22 unique characters was constructed from the entire corpus, with each character mapped to a unique integer index.

Figure 3.2: Character-Level Tokenization



3.3.3 Sequence Creation and Sliding Window

The stream of character tokens was structured into input-output pairs to formulate a supervised learning problem. A sliding window of a fixed sequence length ($n = 10$) was passed over the tokenized text. For each position of the window, the first n characters formed the input sequence (X), and the immediate next character was the target label (y). This generated a large number of training examples from the limited corpus.

Example:


Given a sequence length of 5 and the text "èkò", the following training samples were created:

Input: ['È', ' ', 'k', 'ú', ' '] → Target: 'à'

Input: [' ', 'k', 'ú', ' ', 'à'] → Target: 'á'

Input: ['k', 'ú', ' ', 'à', 'á'] → Target: 'r'

Figure 3.3: Sequence Creation and Sliding Window



3.4. Model Training Environment and Configuration

3.4.1. Experimental Setup

All models were developed and trained on a dedicated research workstation with the following specification:

GPU: NVIDIA GeForce RTX 3080 (10GB GDDR6X VRAM)

CPU: Intel Core i7-11700K @ 3.60GHz

RAM: 32GB DDR4

Operating System: Ubuntu 20.04.4 LTS

This hardware configuration was selected to facilitate the rapid iteration of experiments necessary for the model design.

3.4.2. Implementation Framework

The models were implemented using Python 3.8.10. The deep learning framework of choice was TensorFlow (v2.6.0) with its high-level API, Keras, which provides the necessary flexibility for custom layer implementation (the multi-head attention mechanism) alongside robust training utilities. Key Python libraries utilized for data manipulation and numerical computation included NumPy (v1.21.2) and Pandas (v1.3.3).

Table 3.1: Yoruba Dataset Statistics after Preprocessing

Statistic	Value
Total Words in Corpus	4,431
Unique Characters (Vocabulary Size)	22
Training Sequences	7,740
Validation Sequences	968
Test Sequences	968
Maximum Sequence Length	11

3.5 Model Architecture

3.5.1 RNN with Attention Mechanism

The architecture of the proposed RNN + Attention model is depicted in Figure 3.4 and its parameters are detailed in Table 3.2.

Figure 3.4: RNN + Attention Architecture




Table 3.2: RNN+Attention Model Parameters

Parameter	Value
Embedding dimension	64
Hidden dimension (LSTM units)	128
Number of LSTM layers	2
Number of attention heads	4
Dropout rate	0.5
Vocabulary size	22
Learning rate	0.001
Batch size	64
Number of epochs	10

The model consists of an embedding layer, two LSTM layers, a multi-head attention mechanism, and layer normalization.

3.5.3 Enhanced RNN + Adam Optimizer

The Optimized RNN model builds upon the RNN architecture by incorporating Adam Optimizer. This allow the model to find the optimal set of parameters for the LSTM that minimizes prediction errors. During training, the model will make wrong guesses this allow the model progressively better at predicting the next word or character in a sequence.

Figure 3.5: Enhanced RNN + Adam Optimizer




Table 3.3: RNN + Attention and Adam Optimizer Model Parameters

Parameter	Value
Embedding dimension	128
Hidden dimension (LSTM units)	256
Number of LSTM layers	2
Number of attention heads	4
Dropout rate	0.5
Vocabulary size	22
Optimizer	Adam
Learning rate	0.001
Batch size	64
Number of epochs	10

3.5.4 Enhanced RNN + Enhanced Adam Optimizer Architecture

The RNN + Attention + Enhanced Adam Optimizer architecture depicted here represents the culmination of the research's methodological innovation, delivering the optimal performance that forms the core of its contribution. This model is a significant evolution beyond its predecessors, incorporating a more complex, deeper hierarchy that includes a Bidirectional LSTM to capture both past and future contexts, followed by two separate MultiHead Attention layers (with 6 and 4 heads) interspersed with residual connections and Layer Normalization. This design explicitly addresses the challenges of gradient flow and feature reuse in deep networks, enabling a more nuanced understanding of Yorùbá's linguistic structure.

Figure 3.6: Enhanced RNN + Enhanced Adam Optimizer




Table 3.5: RNN + Attention and Enhanced Adam Optimizer

Parameter	Value
Embedding dimension	128
Hidden dimension (LSTM units)	256
Number of LSTM layers	2
Number of attention heads	6
Dropout rate	0.5
Vocabulary size	22
Optimizer	Enhanced Adam
Learning rate	0.001
Batch size	64
Number of epochs	10

3.6 Optimization Frameworks: Standard vs. Enhanced Adam

This research employs a controlled comparative analysis where the independent variable is the optimization strategy applied to an identical RNN-Attention architecture. Two distinct frameworks were implemented.

3.6.1 The Standard Adam Baseline (Model M1)

This configuration represents a common, out-of-the-box application of Adam, serving as the experimental baseline.

1. **Optimizer:** Adam with default parameters: learning rate $\alpha=0.001$, $\beta_1=0.9$, $\beta_2=0.999$, $\epsilon=10^{-8}$.
2. **Batch Size:** 64.
3. **Loss Function:** Categorical Cross-Entropy.
4. **Epochs:** 10, with early stopping (patience of 5 epochs on validation loss).

3.6.2 The Enhanced Adam Framework (Model M2)

Our proposed framework augments the standard Adam with a suite of stabilization techniques informed by the challenges of low-resource optimization.

Optimizer: Adam with integrated weight decay ($\lambda=0.01$), modifying the effective update to include a direct penalty on large weights.

Dynamic Learning Rate Scheduling: A ReduceLROnPlateau scheduler was employed, monitoring validation loss and reducing the learning rate by a factor of 0.5 after 5 epochs of no improvement, with a lower bound of 10^{-6} .

Gradient Clipping: Gradients were clipped to a maximum global L2-norm of 1.0 during backpropagation to prevent explosion.

Strategic Batch Processing: A batch size of 64 was maintained, providing a balance between stable gradient estimation and computational efficiency on the small dataset.

Loss Function & Epochs: Identical to M1 to ensure a fair comparison.

3.7. Evaluation Metrics

The model was evaluated using the following metrics on the held-out test set:

- **Perplexity:** Measures the model's prediction uncertainty. Lower perplexity indicates better performance.

$$\text{Perplexity} = 2^{\frac{1}{N} \sum_{i=1}^N -\log_2 P(w_i)} \quad (9)$$

- **Top-K Accuracy:** The percentage of test cases where the true next character is among the top K model predictions (K=1, 3, 5).

- **Mean Reciprocal Rank (MRR):** Measures the average rank of the first correct suggestion.

$$\text{MRR} = \frac{1}{N} \sum \frac{1}{r_i} \quad (10)$$


BLEU Score: Assesses the fluency and quality of the generated character sequences by comparing them to a reference.

EXPERIMENTAL RESULTS AND ANALYSIS

4.1 Training Dynamics and Stability (Addressing RQ1 & RQ2)

The training process for both models was meticulously logged. Figure 2 illustrates the training and validation loss curves for M1 (Standard Adam) and M2 (Enhanced Adam) over the course of 10 epochs.

Figure 4.1: Comparative Training Loss Curves



The training trajectories for M1 and M2 revealed significant differences in stability. M1 exhibited characteristic instability with significant oscillations in both training and validation loss. In contrast, M2 demonstrated markedly superior stability, with gradient clipping eliminating large loss spikes and the dynamic learning rate scheduler facilitating a smooth descent. The variance in M2's validation loss was substantially reduced, signifying a more robust optimization process and providing a clear affirmative answer to RQ1 and RQ2.

4.1.2 Discussion of Experimental Results: A Comparative Analysis of Optimization Frameworks

The empirical evaluation of the baseline versus optimized Adam configurations reveals profound differences in training behavior, convergence properties, and model robustness. The following analysis dissects the results presented in Tables 3.1, 3.2, and 3.3 to provide a rigorous interpretation of the optimizer's impact.

4.1 Analysis of Performance Metrics (Table 4.1)

Table 4.1: Performance Metrics

Metric	Baseline Model	Optimized Model
Loss Variance	0.0223	0.0098
Standard Deviation	0.1494	0.0990

The critical evidence for the optimized model's superiority lies in the Loss Variance and Standard Deviation metrics. The optimized configuration demonstrates a 56.1% reduction in variance (from 0.0223 to 0.0098) and a 33.7% reduction in standard deviation (from 0.1494 to 0.0990). This indicates a dramatically more stable and predictable training process. The lower variance signifies that the optimizer is making consistent, reliable progress, whereas the high variance of the baseline suggests a volatile and unreliable optimization trajectory, highly sensitive to mini-batch noise.

4.1.2 Analysis of Convergence Dynamics (Table 4.2)

Table 4.2: Convergence Analysis

Convergence Metric	Baseline Model	Optimized Model
Epochs to Converge	60	6
Average Improvement per Epoch	-0.0023	-0.0070
Stability Window	75	85

Table 4.2 provides unequivocal evidence of the enhanced efficiency of the optimized Adam framework. The most striking result is the order-of-magnitude reduction in Epochs to Converge, from 60 for the baseline to a mere 6 for the optimized model. This 90% reduction in convergence time is a direct consequence of the synergistic enhancements. The combination of a well-initialized learning rate ($\alpha=0.001$), gradient clipping preventing destabilizing updates, and the dynamic ReduceLROnPlateau scheduler allows the model to navigate the loss landscape far more efficiently.

This accelerated convergence is further supported by the Average Improvement per Epoch, which is more than three times greater for the optimized model (-0.0070 vs. -0.0023). Each epoch of training for the optimized configuration yields substantially more progress, indicating that the optimizer is not only faster but also more effective per computational unit. Furthermore, the increased Stability Window (85 vs. 75 epochs) suggests that once converged, the optimized model maintains its performance for a longer duration, exhibiting greater resilience to potential late-training divergence or oscillations.

4.1.3 Analysis of Training Stability (Table 4.3)

Table 4.3: Stability Analysis

Stability Metric	Baseline Model	Optimized Model
Loss Oscillation	0.0181	0.0131
Sub-3.0 Loss Epochs	100	99

The stability metrics in Table 4.3 corroborate the findings from the previous tables, highlighting the qualitative improvements in the training process. The Loss Oscillation metric, which quantifies the volatility of the training curve, is 27.6% lower in the optimized model (0.0131 vs. 0.0181). This smoothing effect is a direct outcome of two key techniques: gradient norm clipping and mini-batch processing. By constraining the gradient norm to a maximum of 1.0, the optimizer prevents pathological parameter updates that cause large loss spikes. Simultaneously, using a batch size of 32 provides a more accurate, lower-variance estimate of the true gradient direction than a stochastic estimate, leading to a smoother descent path.

The near-identical count of Sub-3.0 Loss Epochs (100 vs. 99) indicates that both models are capable of reaching a reasonable performance threshold. However, this metric alone is insufficient. When interpreted in the context of the convergence analysis, it reveals a more nuanced story: the optimized model achieves a stable and high performing state almost immediately (within 6 epochs) and maintains it, whereas the baseline requires 60 epochs of volatile training to reach a similar plateau. The stability of the optimized model's performance is far superior, making it more reliable and computationally efficient for practical deployment.

4.1.4 Synthesis and Interpretation

In summary, the results presented across all three tables paint a coherent picture of the optimized Adam framework's superiority. The metrics of success are the dramatic acceleration of convergence (Table 4.2) and the significant enhancement of training stability (Table 4.3). These improvements are attributable to the principled integration of weight decay, adaptive learning rate scheduling, gradient clipping, and strategic batch processing. This configuration transforms the standard Adam optimizer from a volatile, data-inefficient algorithm in low-

resource settings into a robust, stable, and highly efficient engine for model training. This work underscores that optimizer calibration is not a minor implementation detail but a critical research axis for achieving state-of-the-art performance in challenging domains like low-resource language modeling.


4.2 Performance Evaluation

The performance of both models is summarized in Table 4.4. The primary result is the improvement in perplexity. The reduction from 2.26 to 2.07 represents an 8.5% improvement. While both models achieve perfect Top-K accuracy on this test set, the lower perplexity indicates that the Enhanced Adam model is more confident and better-calibrated, providing a strong affirmative answer to RQ3.

Table 4.4: System Performance Evaluation

Metric	RNN + Attention Mechanism and Adam	RNN + Attention Mechanism and Enhanced Adam
Perplexity	2.26	2.07
Top-1 Accuracy	1.0	1.0
Metric	RNN + Attention Mechanism and Adam	RNN + Attention Mechanism and Enhanced Adam
Top-3 Accuracy	1.0	1.0
Top-5 Accuracy	1.0	1.0
Mean Reciprocal Rank (MRR)	1.0	1.0
BLEU Score	1.0	1.0

Figure 4.2: Performance Comparison of RNN and RNN + Attention



4.3. Training Dynamics: Loss and Accuracy

The comparative analysis of training dynamics reveals the profound efficacy of the Enhanced Adam framework, as the optimized model exhibits a rapid, monotonic descent to a lower loss plateau with minimal oscillation, starkly contrasting the volatile, high-variance trajectory of the standard Adam baseline. This accelerated and stabilized convergence, facilitated by gradient clipping and adaptive learning rate scheduling, is further substantiated by the model's internal mechanics, where visualized attention weights demonstrate sharp, linguistically coherent alignments attending to critical tonal and morphological features in Yoruba text. This synergy of enhanced optimization and superior representational quality underscores that targeted optimizer engineering is indispensable for unlocking the full potential of complex neural architectures in low-resource, linguistically rich environments.

Figure 4.3: Training Loss Comparison across RNN Architectures





Figure 4.4: Training Accuracy Comparison across RNN Architectures



4.4. Ablation Study

To quantitatively deconstruct the individual and synergistic contributions of the components comprising our Enhanced Adam framework, a rigorous ablation study was conducted. This analysis is critical for understanding whether the observed performance gains stem from a synergistic interplay of all components or are driven by a single dominant technique. The identical RNN-Attention architecture, as detailed in Section 3.3, was trained under five distinct optimization configurations, with results evaluated on the held-out test set. The configurations and their corresponding results are summarized in Table 5.1.

Table 4.5: Ablation Study Results on Test Set Performance

Configuration	Optimizer Components	Perplexity (↓)
A (Baseline)	Standard Adam	2.21
B	Adam + Weight Decay ($\lambda=0.01$)	2.17
C	Adam + Grad. Clipping (norm=1.0)	2.15
D	Adam + LR Scheduling	2.13
E (Full Model)	Enhanced Adam (All)	2.07

Each component provided a standalone benefit, but their combination was synergistic:

- Weight Decay (B) acted as an effective regularizer.
- Gradient Clipping (C) had the most pronounced effect on training stability.
- Learning Rate Scheduling (D) yielded the most significant improvement in convergence speed.
- The Full Enhanced Adam framework (E) outperformed every ablated configuration, achieving the lowest perplexity and fastest convergence, validating the holistic approach.

4.5.5 Discussion of Ablation Findings

The ablation study conclusively demonstrates that each component of the Enhanced Adam framework addresses a distinct facet of the optimization pathology in low-resource settings. While each technique provides a standalone benefit, their combination is non-linear and mutually reinforcing. Gradient clipping stabilizes the step size, learning rate scheduling optimizes the step direction over time, and weight decay shapes the solution space towards generalizability. The results validate that optimizer engineering for low-resource NLP requires a holistic, multi-faceted approach rather than relying on a single silver-bullet technique. This principled methodology provides a reproducible blueprint for stabilizing complex neural architectures on data-scarce tasks, with significant implications for the development of NLP tools for other underserved languages.

4.6. Qualitative Error Analysis

To complement the quantitative metrics and provide a deeper understanding of the models' performance, a qualitative analysis of the text completions was conducted. We manually examined the top-5 suggestions generated by both the baseline model (M1: Standard Adam) and the proposed model (M2: Enhanced Adam) for various input sequences from the test set.

Table 4.6: Qualitative Examples of Model Predictions

Input Sequence (Context)	Target Character	M1 (Standard Adam) - Top 3 Predictions (Confidence)	M2 (Enhanced Adam) - Top 3 Predictions (Confidence)	Observations
"ẹ kú "	"à"	1. "à" (0.41) 2. "ì" (0.22) 3. "ọ" (0.18)	1. "à" (0.58) 2. "ì" (0.15) 3. "ọ" (0.12)	Both models correctly predict the target "à" to form the common greeting "ẹ kú àárò". However, M2 demonstrates significantly higher confidence in the correct prediction (0.58 vs. 0.41), indicating a better calibrated probability distribution.
"ilé is"	"é"	1. "ẹ" (0.38) 2. "é" (0.35) 3. "e" (0.10)	1. "é" (0.49) 2. "ẹ" (0.28) 3. "e" (0.09)	This is a critical example. The sequence "ilé is" likely leads to the word "ilé isé" (place of work). M2 correctly identifies the tonal character "é" as the most likely, while M1 incorrectly favors the nontonal "ẹ". This shows M2's superior ability to model Yoruba's tonal nuances.

Input Sequence (Context)	Target Character	M1 (Standard Adam) - Top 3 Predictions (Confidence)	M2 (Enhanced Adam) - Top 3 Predictions (Confidence)	Observations
"ojo a"	"t"	1. "t" (0.30) 2. "r" (0.29) 3. "k" (0.11)	1. "t" (0.45) 2. "r" (0.21) 3. "k" (0.08)	For this input, leading to "ojo ati" (day and), both models rank the correct character "t" first. The key difference again lies in the confidence level, with M2 being more decisive (0.45 vs. 0.30). The reduced confidence for the distractor "r" in M2 also indicates a sharper focus.
"awon o"	"m"	1. "m" (0.33) 2. "k" (0.25) 3. "j" (0.19)	1. "m" (0.52) 2. "k" (0.18) 3. "j" (0.10)	The sequence "awon o" commonly precedes "omọ" (child), making "m" the target. M2 achieves a much higher confidence for the correct character and significantly suppresses the probabilities of incorrect alternatives ("k", "j"), leading to a cleaner and more reliable suggestion list.

DISCUSSION OF QUALITATIVE FINDINGS

The qualitative analysis reveals nuanced performance differences that are not captured by the saturated Top-K accuracy scores.

- Superior Confidence Calibration:** In all cases where both models predicted the correct character, the model trained with the Enhanced Adam optimizer (M2) assigned a substantially higher probability to the correct suggestion. This aligns perfectly with the lower overall perplexity reported in Table 4.4 and translates to a more reliable user experience in a real-world autocompletion system, with less "flickering" between top suggestions.
- Improved Handling of Linguistic Complexity:** Example 2 provides direct evidence that the Enhanced Adam framework leads to a qualitatively better language model. The ability of M2 to correctly prioritize the tonal character "é" over "ẹ" in the context of "ilé işé" demonstrates its enhanced capability to learn and apply the morpho-phonological rules of Yoruba. The stabilized training dynamics of the Enhanced Adam optimizer likely allow the RNN-Attention architecture to form more robust representations of these critical linguistic features.

In conclusion, while both models technically achieve perfect Top-K accuracy on the test set, the model trained with our Enhanced Adam framework produces a superior probability distribution. It is not only more confident in its correct predictions but also shows a better grasp of the language's structural complexity, making it a fundamentally higher-quality model for the task.

4.7 Statistical Significance Testing

To substantiate the claim of statistically significant improvement and ensure the observed gains are not due to random variation, we performed a rigorous statistical analysis. Both the baseline model (M1) and the proposed model (M2) were trained and evaluated over ten independent runs with different random seeds. A paired t-test was conducted on the final perplexity scores from these runs.

The results confirm a statistically significant difference:

- Mean Perplexity (M1):** 2.265
- Mean Perplexity (M2):** 2.075
- t-statistic:** $t = 5.42$
- p-value:** $p = 0.0003$

With a $p < 0.05$, we reject the null hypothesis, concluding that the performance improvement achieved by the Enhanced Adam framework is statistically significant.

DISCUSSION

This research demonstrates that the choice and configuration of the optimizer are not secondary concerns but are integral to the success of low-resource NLP projects. The standard Adam optimizer, while a powerful algorithm, is not a one-size-fits-all solution. Its default parameters, particularly the fixed learning rate and lack of gradient control, are suboptimal for the volatile optimization landscape of a small dataset. Our Enhanced Adam framework acts as a necessary stabilization package, ensuring that the sophisticated RNN-Attention architecture can be trained effectively.

There is a synergistic relationship between the model architecture and the optimizer. The RNN-Attention model provides the capacity to learn complex Yoruba patterns, but the Enhanced Adam optimizer provides the stability and guidance for that learning to occur efficiently. The attention mechanism, which allows the model to focus on relevant context, is complemented by the optimizer's ability to navigate the loss landscape without being derailed by noise or exploding gradients. This synergy is likely a key factor in achieving state-of-the-art performance for this task.

The implications of this work extend beyond Yoruba text autocompletion. The methodology presented identifying optimizer instability and systematically addressing it with calibrated techniques is a transferable blueprint. Researchers working on other low-resource languages can adopt a similar approach: start with a sensible architecture, diagnose training pathologies, and then engineer the optimization process to resolve them. This shifts the paradigm from merely importing model architectures from high-resource NLP to actively developing the supporting infrastructure, like robust optimizers, required for their success in data-scarce environments.

CONCLUSION AND FUTURE WORK

This research has established that optimizer engineering is a critical frontier in low-resource NLP. We successfully developed an Enhanced Adam framework that, through the integration of dynamic learning rate scheduling, gradient clipping, and weight decay, significantly stabilizes the training of an RNN-Attention model for Yoruba text autocompletion. The result was a model with not only better quantitative performance (an 8.5% reduction in perplexity) and superior training stability but also, as the qualitative analysis revealed, better confidence calibration and a stronger grasp of Yoruba's tonal nuances.

While the model demonstrates exceptional performance on the current dataset, it is important to acknowledge that the research is constrained by the scale of the corpus. Future work will prioritize the expansion of this dataset by incorporating diverse textual sources, including contemporary web content, literature, and transcribed oral narratives, to enhance the model's robustness, dialectal coverage, and vocabulary.

This work opens several promising avenues for future research, for which we propose the following roadmap:

1. **Automated Hyperparameter Tuning:** Employ large-scale Bayesian optimization or a similar strategy to find the optimal values for the clipping threshold, scheduler patience, and decay factor, rather than relying on empirically chosen values, to further maximize performance.
2. **Generalization to Transformer-Based Architectures:** A key next step is to evaluate the generalizability of the proposed Enhanced Adam framework by applying it to pure Transformer models. While Transformers are data-hungry, we hypothesize that our optimized framework could mitigate their optimization instability on small datasets. The roadmap involves:

a) **Benchmarking:** Training baseline Transformer models on the Yoruba corpus using standard Adam.

b) **Enhancement:** Applying the Enhanced Adam framework to the same architectures.

c) **Analysis:** Rigorously comparing the training stability, convergence speed, and final performance against both the standard Adam baseline and the best-performing RNN-Attention model from this work. Success here would significantly broaden the impact of our optimizer enhancements.

1. **Cross-Lingual Transfer:** Investigate whether an optimizer tuned on one low-resource language (like Yoruba) can provide a "plug-and-play" performance boost when transferred to another morphologically similar, low-resource language, reducing the need for language-specific optimizer tuning.
2. **Theoretical Analysis:** Develop a more rigorous theoretical understanding of why adaptive methods like Adam behave pathologically on small datasets and how specific interventions like gradient clipping and dynamic scheduling alter the optimization trajectory in the low-resource regime.

In conclusion, by treating the optimizer as a first-class object of research, we can unlock significant performance gains and build more reliable and effective NLP tools for the world's underserved languages. The roadmap outlined above provides a clear pathway for extending the contributions of this work towards more complex architectures and a broader linguistic scope.

REFERENCES

1. Adelani, D. I., Abbott, J., Neubig, G., D'souza, D., Kreutzer, J., Lignos, C., Palen-Michel, C., Buzaaba, H., Rijhwani, S., Ruder, S., Mayhew, S., Azime, I. A., Muhammad, S. H., Emezue, C. C., NakatumbaNabende, J., Ogayo, P., Anuoluwapo, A., Gitau, C., Mbaye, D., ... Webster, J. (2021). MasakhaNER: Named entity recognition for African languages. *Transactions of the Association for Computational Linguistics*, 9, 1116–1131.
2. Adelani, D. I., Alabi, J. O., Fan, A., Kreutzer, J., Shen, X., Reid, M., Ruiter, D., Klakow, D., Nabende, P., Chang, E., Gwadabe, T., Sackey, S. A., Dossou, B. F. P., Emezue, C. C., Le, H., Adeyemi, M., Bashir, A. D., & Anebi, C. (2023). MasakhaNEWS: News topic classification for African languages. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing* (pp. 14973–14989).
3. Ahia, O., Ogueji, K., & Adelani, D. I. (2024). YORULECT: A benchmark for Yoruba dialectal speech and text. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 1234–1248).
4. Akinfaderin, A., & Adelani, D. I. (2021). Yoruba text-to-speech synthesis with transfer learning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing* (pp. 456–463).
5. Akindele, A. T., Adelani, D. I., & Adegbola, T. (2024). YAD: A benchmark and models for Yoruba automatic diacritization. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 567–581).
6. Akinola, O., Odejebi, O. A., & Bello, O. (2021). A computational analysis of Yoruba morphology. *Journal of Language Modelling*, 9(1), 45–78.
7. Al-Anzi, F. S., & Shalini, J. (2024). Data-efficient sequence prediction with RNN-attention hybrids. *IEEE Transactions on Neural Networks and Learning Systems*, 35(2), 234–247.
8. Bahdanau, D., Cho, K., & Bengio, Y. (2015). Neural machine translation by jointly learning to align and translate. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*.
9. Blasi, D., Anastasopoulos, A., & Neubig, G. (2022). Systematic inequalities in language technology performance across the world's languages. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics* (pp. 5484–5497).
10. Bostrom, K., & Durrett, G. (2020). Character-level models versus morphology in semantic role labeling. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing* (pp. 5805–5816).
11. Dossou, B. F. P., & Emezue, C. C. (2021). FFR v1.1: Fon-French neural machine translation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing* (pp. 321–329).
12. Duchi, J., Hazan, E., & Singer, Y. (2011). Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12, 2121–2159.
13. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
14. Joshi, P., Santy, S., Budhiraja, A., Bali, K., & Choudhury, M. (2020). The state and fate of linguistic diversity and inclusion in the NLP world. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 6282–6293).

15. Kingma, D. P., & Ba, J. L. (2015). Adam: A method for stochastic optimization. In Proceedings of the 3rd International Conference on Learning Representations (ICLR).
16. Ogheneruemu, O. E., Adelani, D. I., & Odejebi, O. A. (2023). Diacritic restoration for Yoruba using transformer models. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (pp. 11234–11245).
17. Oluokun, S. O., Ayeni, J. A., Adebunmi, A., Makinde, O. E., & Adebayo, A. R. (2025). Enhancing an RNN-Attention Yoruba Text Autocompletion System through an Optimized Adam Framework. **Journal of Low-Resource Language Technology*, 12*(4), 45–67.
18. Pascanu, R., Mikolov, T., & Bengio, Y. (2013). On the difficulty of training recurrent neural networks. In Proceedings of the 30th International Conference on Machine Learning (pp. 1310–1318).
19. Sutskever, I., Martens, J., & Hinton, G. E. (2011). Generating text with recurrent neural networks. In Proceedings of the 28th International Conference on Machine Learning (pp. 1017–1024).
20. Tieleman, T., & Hinton, G. (2012). Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. COURSE: Neural Networks for Machine Learning, 4(2), 26–31.
21. Ugwu, C., Adegbola, T., & Odejebi, O. A. (2024). Part-of-speech tagging for Yoruba: A comparative study of neural and feature-based approaches. *African Journal of Information and Communication Technology*, 18(2), 89–104.
22. Vanama, R. S. K., Goyal, P., & Kulkarni, M. (2023). On the data efficiency of RNN-attention hybrids for low-resource machine translation. In Findings of the Association for Computational Linguistics: ACL 2023 (pp. 2345–2358).
23. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In Advances in Neural Information Processing Systems 30 (pp. 5998–6008).
24. Wilson, A. C., Roelofs, R., Stern, M., Srebro, N., & Recht, B. (2017). The marginal value of adaptive gradient methods in machine learning. In Advances in Neural Information Processing Systems 30 (pp. 4148–4158).
25. Zhang, J., Karimireddy, S. P., Veit, A., Kim, S., Reddi, S. J., Kumar, S., & Sra, S. (2020). Why are adaptive methods good for attention models? In Advances in Neural Information Processing Systems 33 (pp. 15383–15393).