

NADIR: Differential Attention Flow for Non-Autoregressive Transliteration in Indic Languages

Lakshya Tomar¹, Vinayak Abrol², Puneet Agarwal¹

¹RocketFrog AI

²CSE Department, Indraprastha Institute of Information Technology, Delhi
lakshya.tomar@rocketfrog.ai, abrol@iiitd.ac.in, puneet@rocketfrog.ai

Abstract

In this work, we argue that not all sequence-to-sequence tasks require the strong inductive biases of autoregressive (AR) models. Tasks like multilingual transliteration, code refactoring, grammatical correction or text normalization often rely on local dependencies where the full modeling capacity of AR models can be overkill, creating a trade-off between their high accuracy and high inference latency. While non-autoregressive (NAR) models offer speed, they typically suffer from hallucinations and poor length control. To explore this trade-off, we focus on the multilingual transliteration task in Indic languages and introduce **NADIR**, a novel NAR architecture designed to strike a balance between speed and accuracy. **NADIR** integrates a Differential Transformer and a Mixture-of-Experts mechanism, enabling it to robustly model complex character mappings without sequential dependencies. **NADIR** achieves over a $13\times$ speed-up compared to the state-of-the-art AR baseline. It maintains a competitive mean Character Error Rate of **15.78%**, compared to **14.44%** for the AR model and **21.88%** for a standard NAR equivalent. **Importantly**, **NADIR** reduces Repetition errors by **49.53%**, Substitution errors by **24.45%**, Omission errors by **32.92%** and Insertion errors by **16.87%**. This work provides a practical blueprint for building fast and reliable NAR systems, effectively bridging the gap between AR accuracy and the demands of real-time, large-scale deployment.

Introduction

Transformers (Vaswani et al. 2017) suffer from attention noise (Ye et al. 2025), assigning irrelevant attention scores to unimportant context, which can degrade model performance. This issue is further amplified in Non-Autoregressive (NAR) (Gu et al. 2018) settings due to reduced contextual information, making it difficult for NAR models to match the performance of their Autoregressive (AR) (Sutskever, Vinyals, and Le 2014) counterparts. This gap is largely due to insertions, substitutions, omissions, and repetitions errors, which we collectively refer to as **NAR Hallucinations**. However, we argue that not all tasks benefit from the strong sequential bias of AR models. Tasks such as multilingual transliteration, code refactoring, and grammatical correction often rely more on local dependencies. In this

work, we show that reducing attention noise and incorporating a Mixture-of-Experts (MoE) (Shazeer et al. 2017) architecture enables NAR models to achieve competitive performance while being significantly more efficient. To validate our approach, we focus on Indic transliteration as a representative task.

Indic scripts include a diverse family of writing systems such as Devanagari (used in Hindi, Marathi, Sanskrit), Bengali, and Punjabi/Gurmukhi, which have been deeply rooted in South Asian cultures for centuries. Today, more than **1.6 billion people** use these scripts in various languages. A key linguistic challenge with these scripts is transliteration, the process of converting text from one script to another while preserving pronunciation. Transliteration differs from translation: It maps the sounds of a word, not its meaning. Transliteration is inherently challenging due to: (a) ambiguity in character mappings, including many-to-one, one-to-many, and many-to-many relationships, (b) phonetic variability, where different words in a non-English language may be transliterated into the same Roman word, and (c) homophones and phonological constraints, where similar sounds are represented by different characters depending on context. Thus, the context of surrounding characters is essential for accurate transliteration, which is why AR models typically deliver superior performance.

While there is a growing shift from AR to NAR models in related tasks, techniques such as Knowledge Distillation (Gong, Zhou, and Qian 2022), Iterative Refinement (Lee, Shu, and Cho 2020), and CTC Loss (Yu et al. 2025), as detailed in Related Work, have been employed to mitigate the accuracy loss incurred from abandoning AR approaches. However, to our knowledge, this issue remains unresolved for transliteration tasks. In this work, to address this need, we propose **NADIR (Non-Autoregressive Differential Intelligent Router)**, a multilingual deep learning architecture inspired by the differential transformer and Mixture-Of-Expert, capable of transliterating **180.1k** words in just **2 minutes and 55 sec.** (≈ 1005 words/second). In contrast, the current state-of-the-art IndicXLIT model (Madhani et al. 2023) requires **38 minutes and 50 sec.** (≈ 77 words/second). Also, Our model achieves a competitive mean CER of **15.78%** (STD: 5.67%) versus IndicXLIT’s **14.44%** (STD: 4.35%). The key contributions of our work are as follows:

1. We present an analysis of failure modes that arise when

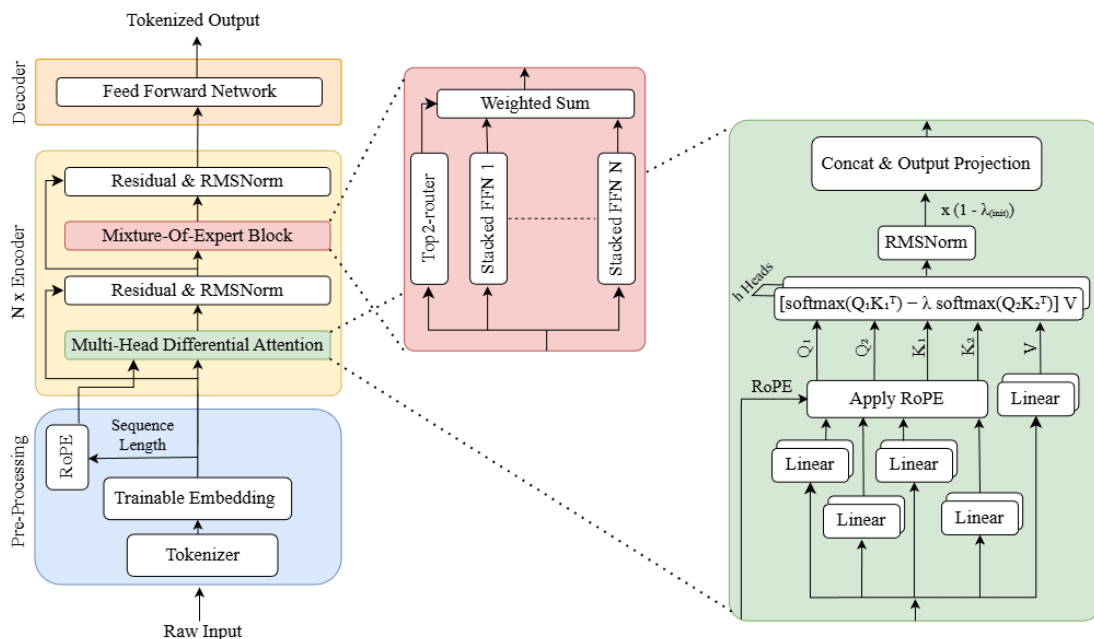


Figure 1: *NADIR*: Architecture Overview showing pre-processing, multiple encoder and non-autoregressive decoder blocks. On the right, we show an expanded view of the encoder block that is built upon differential transformer layers combined with MoE-based routing, enabling efficient and expressive context modelling.

switching from autoregressive to non-autoregressive models, collectively referred to as *NAR Hallucinations*.

2. **First Non-Autoregressive Transliteration:** State-of-the-art transliteration models have very slow decoding speed, making them impractical for large-scale or latency-sensitive applications. We address this gap by introducing *NADIR* a novel transliteration model designed for extreme throughput without sacrificing accuracy.

- (a) **Scalable Transliteration via *NADIR*:** To the best of our knowledge, this is the first work to address the hallucination problem in non-autoregressive NLP models for transliteration.
- (b) Our approach integrates Differential Attention Mechanism (Ye et al. 2025), which significantly reduces repetition, substitution and omission errors by minimizing attention noise. Complementing this, the MoE module further reduces insertion and remaining repetition and substitution errors by enabling dynamic, token-specific computation — as confirmed by our ablation study.

In summary, the research question we addressed: Can reducing attention noise and incorporating MoE help NAR models capture context effectively without auto-regression?

Next, we position our work in the context of prior research in the Related Work section. We then present a detailed description of the proposed *NADIR* architecture, along with the underlying intuition behind its design. This is followed by a brief overview of the datasets used and a comprehensive explanation of the training setup, including model parameters and optimization strategies. In the Experimental Results sec-

tion, we report empirical comparisons against strong AR and NAR baselines, highlighting *NADIR*'s performance across multiple error categories. We also include an ablation analysis to assess the individual contributions of key architectural components. Finally, we conclude with a summary of our findings and outline directions for future work.

Related Work

NAR models have emerged as a promising alternative to traditional AR methods for sequence-to-sequence tasks, offering significant inference speed-ups. However, this parallelism often comes at the cost of ‘hallucinations’ errors such as token repetitions, substitutions, omissions and insertions, due to the model’s inherent conditional independence assumption. Previous research has explored several strategies to mitigate these issues, each with its own set of trade-offs.

An early and prominent line of work involves leveraging an AR ‘teacher’ model to guide the training of an NAR ‘student’ through knowledge distillation (Gong, Zhou, and Qian 2022). While such methods have demonstrated improved performance, this dependency on a pre-trained AR model introduces significant complexity, increases overall training time, and moves away from a truly end-to-end parallel training paradigm. Other approaches have focused on modifying the loss function or model architecture. The Connectionist Temporal Classification (CTC) loss, particularly explored in tasks like speech recognition and Chinese-Braille translation (Yu et al. 2025), facilitates alignment but is limited by its assumption of monotonic alignment and conditional independence between output tokens. These assumptions are often violated in complex transliteration and translation sce-

narios. Similarly, early architectures like FlowSeq (Ma et al. 2019) attempted to address NAR limitations by incorporating a small auxiliary network for explicit length prediction. However, an inaccurate length prediction can trigger cascading errors, destabilizing the entire generation process.

To address the core issue of conditional independence, NAR paradigms have evolved, primarily in the field of machine translation with the most successful approaches being iterative refinement. Instead of generating the entire sequence in a single pass, these models iteratively edit a draft translation over a fixed number of steps. Seminal models like Mask-Predict (Ghazvininejad et al. 2019) use a conditional masked language model objective to re-predict subsets of tokens that it is least confident about. Similarly, the Levenshtein Transformer (Gu, Wang, and Junbo 2019) explicitly trains the model to perform editing operations like insertion and deletion. While these methods significantly close the quality gap with AR models, they trade some of the latency gains for higher accuracy. Another advanced technique involves incorporating latent variables to guide the parallel decoding process. Models like the Disentangled Context Transformer (Kasai et al. 2020) and Glancing Transformer (Qian et al. 2021) learn latent alignments or fertilities that provide a soft structural blueprint for the output, improving coherence without full sequential conditioning. A more recent work has also explored energy-based models (Tu et al. 2020) to improve the quality of the generated sequences.

Despite these advancements in general sequence-to-sequence modeling, the specific application of advanced NAR techniques to multilingual transliteration remains underexplored. Transliteration presents a unique challenge: it requires strict phonetic and orthographic accuracy but is more constrained and localized than open-domain translation. The overhead of iterative refinement may be unnecessary, while the simple conditional independence of basic NAR models is insufficient. This highlights a clear need for an architecture that can achieve a robust balance—capturing local dependencies with high fidelity while retaining maximum parallelism. Our work, *NADIR*, is designed to fill this gap by integrating a differential attention mechanism and expert Gating Network (Shazeer et al. 2017), tailored specifically for the demands of high-throughput, multilingual transliteration. In contrast to existing works, the proposed NAR model enables fully parallel and efficient training/decoding. We demonstrate that token-level cross-entropy loss is enough to train the model & this avoids CTC’s rigid alignment and independence assumptions, leading to better fluency and reordering capability. In addition, *NADIR* bypasses explicit length prediction, removing a major source of instability in prior NAR systems.

NADIR (Non-Autoregressive Differential Intelligent Router)

We introduce *NADIR*, a NAR multilingual transliteration model designed to address hallucination problems in transliteration, while maintaining the speed and competitive performance compared to SOTA AR models. As illustrated

in Figure 1, the overall pipeline begins with a preprocessing stage that applies a tokenizer along with learnable token and rotary positional embeddings (RoPE) (Su et al. 2024). Our core innovation lies in the stacked encoder blocks, each built upon Differential Transformer (Ye et al. 2025) layers combined with MoE-based routing, enabling efficient and expressive context modeling. A lightweight MLP-based non-autoregressive decoder generates output characters (to the target script vocabulary) in parallel using the refined encoder representations.

Multi-head Differential Attention: In NAR, the absence of the sequential inductive bias makes it difficult for the standard attention mechanism to focus on the most relevant input tokens. This often leads to noisy attention maps, because of which we observed issues like token insertions, substitutions, omissions, or repetitions. To mitigate this, one can use **Differential Attention Mechanism**, that computes the difference between normalized-softmax attention scores and regulates it using a learnable parameter λ . Formally, given a d -dim input sequence $X \in \mathbb{R}^{M \times d}$ & length M , the query and key projections are each partitioned into two components (Q_1, Q_2 and K_1, K_2) as

$$[Q_1; Q_2] = XW^Q; \quad [K_1; K_2] = XW^K; \quad V = XW^V$$

The attention output is computed by modulating the primary attention score with the secondary one, controlled by a learnable scalar λ :

$$\text{DiffAttn}(X) = \left(S \left(\frac{Q_1 K_1^\top}{\sqrt{d}} \right) - \lambda S \left(\frac{Q_2 K_2^\top}{\sqrt{d}} \right) \right) V \quad (1)$$

Here, $S(\cdot)$ is the softmax function, and W^Q, W^K, W^V are the projection matrices. To ensure stable training, the dynamic modulator λ is parameterized using learnable vectors (λ_q, λ_k) and an initial bias λ_{init} :

$$\lambda = \exp(\lambda_{q_1} \cdot \lambda_{k_1}) - \exp(\lambda_{q_2} \cdot \lambda_{k_2}) + \lambda_{\text{init}}$$

A key yet important detail in *NADIR* is the use of **RM-SNorm** (Zhang and Sennrich 2019), which empirically performs consistently better than the **GroupNorm** (Wu and He 2018) in the conventional differential attention block.

Mixture-of-Expert Module: Error analysis of the differential transformer based model (without MoE) demonstrated better performance on languages that had more training data than others. This observation led us to hypothesize that a single shared Feed-Forward Network (FFN) might not effectively capture the diversity of scripts for all languages. As an initial step, we replaced a single FFN in each layer with a set of small FFNs, referred to as “expert” and manually routed tokens based on linguistic or regional knowledge. While this hardcoded routing showed promising results, particularly for low-resource languages, its effectiveness was likely limited by the small capacity of each expert. One workaround could be scaling expert size based on data availability per cluster, but such heuristics break generality and are not scalable. To address this, we transitioned to a Mixture-of-Experts (MoE) framework, where routing is learned dynamically. A learned router FFN assigns tokens

to experts, which potentially enables flexible and context-aware specialization. This MoE design showcases a bit of robustness in multilingual settings - as shown in the ablation study. Formally MoE layer contains M expert feed-forward networks, denoted as $\{E_1, E_2, \dots, E_M\}$. For each input token representation $x \in \mathbb{R}^d$, a trainable gating network $G(x)$ computes routing probabilities over all experts:

$$p_i = \frac{\exp(g_i(x))}{\sum_{j=1}^M \exp(g_j(x))}, \quad \text{for } i = 1, \dots, M$$

where $g_i(x)$ is the logit score for the i^{th} expert. In our setup, we adopt **Top-2 routing** (Fedus, Zoph, and Shazeer 2022), where only the two experts with the highest gating scores are selected for each input token. Let i and j be the indices of the top-2 experts, then the final output is computed as:

$$\text{MoE}(x) = p_i \cdot E_i(x) + p_j \cdot E_j(x)$$

Implicit Sequence Termination and Training Objective

A fundamental challenge for NAR models is determining when to terminate generation. Unlike AR models, which naturally predict an end-of-sequence ($[\text{EOS}]$) token, NAR models generate all tokens in parallel. While prior work often resorts to a separate, complex length-predictor network, this can introduce instability.

We adopt a more direct approach by enabling the NAR model to perform implicit length prediction. To achieve this, we append an $[\text{EOS}]$ token to every target sequence during training. The key innovation lies in our loss computation: it is calculated only over the token positions up to and including the first predicted $[\text{EOS}]$ token, thereby explicitly encouraging the model to learn when to stop generation. During inference, the model generates tokens up to a fixed maximum length, and we extract the output up to the first position where the model predicts the $[\text{EOS}]$ token. This strategy elegantly compels the model to learn sequence boundaries without an auxiliary network, though its effectiveness is highest in tasks with clearly defined termination points, e.g., it will not generalize to tasks with highly ambiguous or variable-length outputs. To support this mechanism, we employ a composite loss function, which is a weighted sum of two key components, namely;

1. **Token-Level Cross-Entropy Loss** to ensure local prediction accuracy. Let $\hat{\mathbf{Y}} \in \mathbb{R}^{B \times T \times V}$ be the predicted logits for a batch of size B , with sequence length T and vocabulary size V , and $\mathbf{Y} \in \mathbb{N}^{B \times T}$ the ground-truth targets, the loss is computed as

$$\mathcal{L}_{\text{token}} = \frac{1}{B \cdot T} \sum_{b=1}^B \sum_{t=1}^{T_b} -\log P(Y_{b,t} | \hat{Y}_b)$$

2. **Load-Balancing Loss (Shazeer et al. 2017)** to ensure uniform expert utilization in our Mixture-of-Experts (MoE) layers. Let $\mathbf{G} \in \mathbb{R}^{B \times M}$ denote the gating probabilities over M experts for each input in the batch, the loss is computed as:

$$\mathcal{L}_{\text{load}} = M \cdot \sum_{e=1}^M \left(\frac{1}{B} \sum_{b=1}^B G_{b,e} \right)^2$$

The final training objective is a weighted sum of the above:

$$\mathcal{L}_{\text{total}} = \alpha \mathcal{L}_{\text{token}} + \beta \mathcal{L}_{\text{load}}$$

where α , and β are scalar hyperparameters controlling the contribution of each loss component.

Why Differential Transformer Helps in Transliteration

In a standard multi-head attention (Vaswani et al. 2017) layer, the outputs of all heads are aggregated, typically through concatenation followed by a linear projection. However, this aggregation can lead to ambiguous representations. If some heads capture features for the correct output while others capture a simpler, incorrect alternative, the final vector is muddled as a blend of both. The Differential Transformer enhances this process. Instead of a simple sum, it computes a weighted sum over the attention heads, operating on the differences between adjacent head states. This ‘‘differential softmax’’ reintroduces local context and selectivity. This approach offers two key advantages for NAR transliteration. First, it encourages smoother, more stable layer-wise refinement of token representations. The model learns how much of the ‘‘difference’’ to apply at each step, preventing abrupt changes that cause hallucinations. Secondly, it provides fine-grained, content-aware control over how features evolve through the layers. This is crucial for transliteration, where local phonetic consistency is key, as it aligns better with the task’s monotonic structure. For instance, consider the input ‘ksha’ (devnagari word). Additive heads might capture the general ‘complex consonant cluster’ feature, while a subtractive head can learn to specifically identify the most common simplification, like the feature for ‘ka’. The final representation becomes: $[(\text{cluster feature}) - \lambda (\text{‘ka’ feature})]$, actively carving the ‘ka’ ambiguity out of the representation, leaving a sharper, more precise vector that points decisively to ‘ksha’.

Experimental Setup

This section we provide experimental protocol and the dataset used in the experimental study.

Dataset

We evaluate the performance of various transliteration models on the Aksharantar dataset (Madhani et al. 2023), which is the largest open-source parallel dataset available for transliteration in multiple Indian languages. This dataset includes parallel word-level mappings between English and 21 Indic languages and vice versa. Aksharantar dataset includes parallel corpora across 21 Indian languages. The largest training sets are for Malayalam (4.1M), Tamil (3.2M), Kannada (2.9M), and Telugu (2.4M). Languages like Gujarati, Hindi, and Bengali each contribute over 1M training samples. Several mid- to low-resource languages, including Urdu (699k), Konkani (612k), Panjabi (514k), and Oriya (346k), are also represented. Low-resource languages such as Kashmiri (46k), Sindhi (59k), Bodo (35k), Manipuri (106k) included to test generalization. The total dataset comprises **24.8 million** training, **129.6k** validation, and **180.1k** test samples.

Language	Roman → Indic						Indic → Roman					
	CER ↓		WAcc ↑		InfT (sec.) ↓		CER ↓		WAcc ↑		InfT (sec.) ↓	
	IndicXLIT	NADIR	IndicXLIT	NADIR	IndicXLIT	NADIR	IndicXLIT	NADIR	IndicXLIT	NADIR	IndicXLIT	NADIR
Telugu (Tel)	11.03	9.15	58.66	66.92	136.22	10.42	11.97	12.50	41.85	41.43	152.82	10.37
Maithili (Mai)	11.67	12.60	61.24	58.56	165.48	5.34	15.79	15.09	37.24	39.48	64.65	5.62
Bengali (Ben)	15.01	18.06	52.29	46.29	189.94	14.11	20.87	21.34	25.09	24.16	212.87	14.05
Nepali (Nep)	9.93	10.80	65.18	64.76	36.00	4.03	8.31	8.90	65.76	64.74	39.78	4.32
Assamese (Asm)	12.29	16.65	52.8	43.11	59.72	5.23	14.8	15.67	40.27	37.92	65.81	5.80
Malayalam (Mal)	14.29	12.85	43.31	55.39	165.48	12.77	13.76	14.08	33.28	33.3	186.32	12.73
Urdu (Ur)	18.38	21.13	43.58	37.08	202.84	14.94	26.35	27.07	17.58	16.08	226.42	14.62
Hindi (Hin)	11.45	13.47	58.18	54.01	131.26	10.14	15.17	15.80	39.73	38.56	146.69	10.29
Sindhi (Sid/Sin)	19.44	22.15	47.86	38.87	76.90	6.71	23.32	24.12	23.54	22.32	84.36	7.34
Tamil (Tam)	17.7	12.40	28.72	60.95	149.69	11.82	17.77	18.14	30.97	30.52	166.39	11.89
Manipuri (Man)	9.59	14.05	65.43	53.28	47.57	4.46	13.51	18.06	49.5	39.39	52.54	5.13
Sanskrit (San)	15.21	14.88	49.77	51.56	56.56	5.24	11.71	12.50	44.44	43.59	62.56	5.54
Gujarati (Guj)	11.02	12.31	58.94	54.88	251.54	17.97	21.58	21.71	22.01	21.38	285.41	17.70
Punjabi (Pan)	18.47	20.78	43.39	38.42	141.14	11.29	20.78	21.76	30.11	28.07	158.38	11.19
Konkani (Kok)	14.46	15.65	51.63	48.67	51.07	4.98	17.89	19.41	31.48	28.34	56.38	5.33
Kashmiri (Kas)	27.46	34.32	25.51	15.35	86.30	6.88	26.71	29.60	19.3	16.00	96.76	7.14
Boro (Brx)	14.13	16.20	52.17	47.90	35.61	4	13.25	14.91	48.96	44.69	39.10	4.31
Odia (Ori)	17.55	18.05	41.56	39.75	37.61	4.31	12.33	13.50	47.66	44.58	41.67	4.49
Kannada (Kan)	9.37	8.04	62.02	68.36	149.71	11.53	9.68	10.38	44.23	43.03	167.00	11.19
Marathi (Mar)	10.29	12.09	62.46	58.43	158.91	12.21	16.42	16.67	32.99	32.43	176.73	12.4
Mean	14.44	15.78	51.23	50.13	116.48	8.95	16.59	17.56	36.29	34.5	124.18	9.07
Std. Dev.	4.35	5.67	10.87	12.21	62.71	4.14	5.10	5.3	11.75	11.5	71.03	3.95

Table 1: Comprehensive comparison of baseline (IndicXLIT) and proposed model (*NADIR*) for transliteration in both Roman ↔ Indic directions. Metrics include Character Error Rate (CER), Word-level Accuracy (WAcc), and Inference Time (InfT). For each metric, the better result per language is in bold

Model Training Setup

We trained two models, one for each direction, i.e., Roman-to-Indic and Indic-to-Roman. Both models were trained for 100 epochs using 4 encoder layers. Each encoder layer contains 8 attention heads and 5 experts. Each expert comprises a stack of FFN($\text{embed_dim} \times \frac{\text{expert_dim}}{2}$) → GELU → FFN($\frac{\text{expert_dim}}{2} \times \text{expert_dim}$) → GELU → FFN($\text{expert_dim} \times \text{embed_dim}$). Where expert_dim and embed_dim are 512, 768 respectively. We use the AdamW (Loshchilov and Hutter 2019) optimizer with a learning rate of 1×10^{-3} and a weight decay of 1×10^{-3} , along with a linear learning rate scheduler (Wolf et al. 2020) where 15% of the total steps are allocated for warmup, with Dropout (Srivastava et al. 2014) of 0.1 and Capacity Factor (Shazeer et al. 2017) of 1.25. The best performance on *NADIR* was observed when hyperparameters α , and β were set to **0.8**, and **0.2**, in our experiments.

The total number of parameters in *NADIR* is approximately **27 million**. Training is conducted on two NVIDIA RTX 3090 GPUs, inference is performed using a single GPU with a batch size of 8192.

Experimental Results

To evaluate our hypotheses, we present a comparative analysis of performance and inference speed for the proposed *NADIR* model and IndicXLIT, the current state-of-the-art autoregressive model for transliteration, under various experimental scenarios. We systematically vary key components of the model to understand their individual and combined effects on transliteration performance.

Ground Truth	NAR Output	AR Output
direktorrao	direktararao	direktorrao
mushtaidi	mushtaadad	mushtaidi
undannadi	undhanna	undannadi
samvardhita	sambabababadhadhata	samvardhita

Table 2: Examples showing hallucinations in NAR outputs compared to AR outputs. All outputs are in Roman script for readability.

Main Result

In Table 1, we present a detailed comparison between the proposed model *NADIR* and the baseline model IndicXLIT in both transliteration directions (Roman ↔ Indic) across 20 Indian languages. We evaluate the models using three metrics: **Character Error Rate (CER ↓)**, **Word Accuracy (WAcc ↑)**, and **Inference Time (InfT ↓)**.

Indic→Roman direction: On average, across all 20 languages in the Indic→Roman transliteration task, *NADIR* achieves a CER of **17.56**, which is comparable to IndicXLIT’s **16.59**. The average WAcc of *NADIR* is **34.5%**, closely matching IndicXLIT’s **36.29%**. More importantly, *NADIR* offers a substantial efficiency advantage—reducing the mean InfT from **124.18 seconds** to just **9.07 seconds**, representing an order-of-magnitude speedup.

Roman→Indic direction Across all 20 languages, *NADIR* achieves an average CER of **15.78**, closely aligned with IndicXLIT’s **14.44**. The WAcc of *NADIR* is **50.13%**,

Model	Insertion	Substitution	Omissions	Repetition
Standard NAR	28,454	72,127	37,769	6,313
Diff NAR	27,682	57,349	23,453	4,126
Gain over Standard NAR	2.71%	20.49%	37.90%	34.64%
Diff MoE NAR	23,654	54,494	25,334	3,186
Gain over Standard NAR	16.87%	24.45%	32.92%	49.53%
Gain over Diff NAR	14.55%	4.98%	-8.02%	22.78%

Table 3: Hallucination Error Breakdown count across NAR Variants (lower is better). Gains are shown with respect to the baseline (Standard NAR) and the previous row.

Ground Truth	Encoders using Standard Attention	Encoders using Differential Attention
direktorrao	direktararao	direktorrao
mushtaidi	mushtaadad	mushtaidi
undannadi	undhanna	undannadi
mononayonpotro	monoyoyonpot	mononayonpotro
mononoyonprotyaashi	monoyoyprapryaashi	mononoyonprotyaashi

Table 4: Examples showing hallucinations in NAR outputs and how Differential Attention reduces such errors.

slightly below IndicXLIT’s **51.23%**. Once again, *NADIR* demonstrates a significant reduction in inference latency, bringing down the mean InfT from **116.48 seconds** to just **8.95 seconds**. Notably, despite similar overall averages, *NADIR* outperforms IndicXLIT in both CER and WAcc for **5 out of the 20 languages**, underscoring its robustness across diverse scripts.

We also consider the difference between metrics for *IndicXLIT* and *NADIR* models, denoted by Δ . For the Indic→Roman direction, the mean and standard deviation of Δ CER were **-0.96** and **-0.2**, respectively, indicating slightly better character-level accuracy for IndicXLIT. The corresponding values for Δ WAcc were **1.80** and **0.25**. Inference Time, however, shows a substantial gain for *NADIR*, with a mean Δ InfT of **115.11 seconds** and a standard deviation of **67.08**. The Roman→Indic direction follows a similar trend: Δ CER had a mean and standard deviation of **-1.34** and **-1.32**, while Δ WAcc values were **1.1** and **-1.34**. Inference time gains remained consistent, with a mean Δ InfT of **107.53 seconds** and a standard deviation of **58.57**.

These results demonstrate that while *NADIR* slightly trails IndicXLIT in accuracy, it delivers comparable overall performance with a dramatic reduction in inference time, making it a highly practical choice for real-time and large-scale deployment scenarios.

Model	mean CER	mean WAcc
Standard Attention	21.88	38.98
Differential Attention	16.12	46.89
Diff-Attn + MoE	15.78	50.13

Table 5: Average CER and Accuracy across all languages for different model variants.

Ablation Study

In our experiments with the standard encoder-based NAR model, we observed all four types of NAR Hallucination, namely, character insertions, substitutions, omissions, and repetitions. Table 2 presents representative failure cases, highlighting these Hallucinations, by comparing AR and NAR predictions.

We measured the extent of NAR hallucinations across the outputs of three model variants: a) Standard NAR, b) Differential Transformer-based NAR (Diff NAR), and c) Differential Transformer with Mixture-of-Experts (Diff MoE NAR). Insertion, substitution, and omission errors were quantified using the `editdistance` algorithm. Repetition errors were further categorized into three types: Insert Repeat (repeated spans not present in the ground truth), Substitute Repeat (repeated spans that replace expected content but are not part of the ground truth), and Valid Repeat (spans present in the ground truth but repeated more times than necessary). We considered character spans ranging from bigrams to four-grams and counted each distinct repeated span only once, regardless of its frequency. The results of these measurements are summarized in Table 3.

Impact of Differential Attention Mechanism Switching from the standard to the Differential Attention Mechanism helps alleviate the **NAR Hallucination** problem discussed in the previous section. We hypothesize that the subtraction operation in Equation 1 effectively suppresses attention noise, thereby minimizing the influence of irrelevant context that often leads to hallucinations, particularly in the absence of strong autoregressive signals. This hypothesis is supported by the quantitative results presented in Table 3, where the Differential NAR model demonstrates significant improvements in addressing substitutions, omissions, and repetitions, though it shows limited impact on reducing insertions. Qualitative examples illustrating these improve-

ments are provided in Table 4, and they align closely with the statistical trends observed.

Ground Truth	Encoders without MOE	Encoders with MOE
ṭore	ṭor	ṭore
mahama	mahamam	mahama
anukarana	anukaranna	anukarana
emmelayelapai	mlalapai	emmelayelapai
dasakare	dasakarere	dasakare

Table 6: Examples showing how MoE enhances script-specific accuracy over Differential Attention-only models. All text is in Roman script for clarity.

Impact of Mixture-of-Experts (MoE) with Differential Attention Mechanism While Differential Attention significantly mitigates **NAR hallucinations**, models trained solely with this mechanism continue to struggle with the introduction of unwanted characters, particularly insertions (as shown in Table 3). Moreover, there remains substantial room for improvement in other hallucination types.

To address these limitations, we incorporate a Mixture-of-Experts (MoE) module within each encoder layer. This addition yields a notable reduction in insertion errors and also shows meaningful improvements in handling repetitions. While the MoE-NAR variant achieves a modest gain in reducing substitution errors, it does introduce a slight increase in omission errors. Table 6 provides qualitative examples where the MoE-enhanced architecture successfully corrects errors that the Differential Attention Mechanism alone fails to resolve. Additionally, as shown in Table 3, the incorporation of the MoE module leads to a reduction in Insertion, Substitution, and Repetition errors by **14.55%**, **4.98%**, and **22.78%**, respectively albeit with an approximate **8%** increase in omission errors.

Effect of Batch Size on Inference Time Figure 2 illustrates the inference time (in seconds) of *NADIR* and *IndicXLIT* models across varying batch sizes on a log scale. As expected, both models benefit from increased batch sizes up to a point, after which inference time starts to rise due to memory and hardware constraints. *NADIR* consistently outperforms *IndicXLIT* across all batch sizes, demonstrating significantly lower latency and better scalability. Notably, while *IndicXLIT* exhibits a sharp U-shaped curve with a narrow optimal batch window, *NADIR* achieves near-optimal performance over a much wider range, highlighting its robustness and suitability for high-throughput deployment.

Summary

To summarize the findings, transitioning from AR to NAR introduces consistent errors, collectively referred to as NAR Hallucination. The Differential Transformer effectively addresses most of these issues and serves as the primary contributor to performance improvements. Additional enhancements, such as incorporating Mixture-of-Experts (MoE),




Figure 2: Inference Time vs. Batch Size (log scale). *NADIR* achieves significantly lower inference latency compared to the AR-based *IndicXLIT*, especially at large batch sizes.

further help mitigate difficult edge cases where even the Differential Transformer falls short. Table 5 reports the mean CER and WAcc for each model variant.

Conclusion

In this work, we revisit the design space of non-autoregressive models for sequence-to-sequence tasks where local dependencies dominate, using multilingual transliteration as a representative case. We present **NADIR**, a novel NAR architecture that combines a Differential Transformer with a Mixture-of-Experts framework to tackle key challenges such as hallucinations, poor length control, and loss of linguistic fidelity.

Our findings demonstrate that *NADIR* delivers significant improvements over standard NAR baselines while achieving an order-of-magnitude reduction in inference time compared to autoregressive (AR) models. Crucially, *NADIR* reduces repetition, substitution, omission, and insertion errors substantially, thereby closing the accuracy gap with AR models without inheriting their computational overhead.

This work provides both a conceptual and empirical foundation for designing high-throughput, accurate NAR systems that are well-suited for real-time applications in resource-constrained or large-scale multilingual settings. Looking ahead, we believe that the principles underlying *NADIR* differential attention, expert specialization, and hallucination aware evaluation can be extended to a broader class of structured generation tasks beyond transliteration. We also aim to extend *NADIR* to tasks beyond transliteration, with a focus on incorporating recent advancements in Mixture-of-Experts into the architecture.

Acknowledgements

This work was supported by Infosys Foundation via the Infosys Centre for Artificial Intelligence, Indraprastha Institute of Information Technology Delhi, and in part by the Nebius Research Grant.

References

- Fedus, W.; Zoph, B.; and Shazeer, N. 2022. Switch Transformers: Scaling to Trillion Parameter Models with Simple and Efficient Sparsity. *Journal of Machine Learning Research*, 23(120): 1–39.
- Ghazvininejad, M.; Levy, O.; Liu, Y.; and Zettlemoyer, L. 2019. Mask-Predict: Parallel Decoding of Conditional Masked Language Models. In *Conference on Empirical Methods in Natural Language Processing and International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 6112–6121.
- Gong, X.; Zhou, Z.; and Qian, Y. 2022. Knowledge Transfer and Distillation from Autoregressive to Non-Autoregressive Speech Recognition. In *ISCA Interspeech*, 2618–2622.
- Gu, J.; Bradbury, J.; Xiong, C.; Li, V. O.; and Socher, R. 2018. Non-Autoregressive Neural Machine Translation. In *International Conference on Learning Representations (ICLR)*.
- Gu, J.; Wang, C.; and Junbo, J. Z. 2019. Levenshtein transformer. In *Advances on Neural Information Processing Systems (NeurIPS)*, 11181 – 11191.
- Kasai, J.; Cross, J.; Ghazvininejad, M.; and Gu, J. 2020. Non-autoregressive machine translation with disentangled context transformer. In *International Conference on Machine Learning (ICLR)*.
- Lee, J.; Shu, R.; and Cho, K. 2020. Iterative Refinement in the Continuous Space for Non-Autoregressive Neural Machine Translation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1006–1015.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations (ICLR)*.
- Ma, X.; Zhou, C.; Li, X.; Neubig, G.; and Hovy, E. 2019. FlowSeq: Non-Autoregressive Conditional Sequence Generation with Generative Flow. In *Conference on Empirical Methods in Natural Language Processing and International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 4282–4292.
- Madhani, Y.; Parthan, S.; Bedekar, P.; Nc, G.; Khapra, R.; Kunchukuttan, A.; Kumar, P.; and Khapra, M. 2023. Aksharantar: Open Indic-language Transliteration datasets and models for the Next Billion Users. In *Findings of the Association for Computational Linguistics (EMNLP)*, 40–57.
- Qian, L.; Zhou, H.; Bao, Y.; Wang, M.; Qiu, L.; Zhang, W.; Yu, Y.; and Li, L. 2021. Glancing Transformer for Non-Autoregressive Neural Machine Translation. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 1993–2003.
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q.; Hinton, G.; and Dean, J. 2017. Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer. In *International Conference on Learning Representations (ICLR)*.
- Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 15(56): 1929–1958.
- Su, J.; Ahmed, M.; Lu, Y.; Pan, S.; Bo, W.; and Liu, Y. 2024. RoFormer: Enhanced Transformer with Rotary Position Embedding. *Neurocomputing*, 568: 127063.
- Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to Sequence Learning with Neural Networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 3104 – 3112.
- Tu, L.; Pang, R. Y.; Wiseman, S.; and Gimpel, K. 2020. ENGINE: Energy-Based Inference Networks for Non-Autoregressive Machine Translation. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2819–2826.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 6000 – 6010.
- Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; Davison, J.; Shleifer, S.; von Platen, P.; Ma, C.; Jernite, Y.; Plu, J.; Xu, C.; Le Scao, T.; Gugger, S.; Drame, M.; Lhoest, Q.; and Rush, A. M. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Conference on Empirical Methods in Natural Language Processing (EMNLP Demonstrations)*.
- Wu, Y.; and He, K. 2018. Group Normalization. In *European Conference on Computer Vision (ECCV)*, 3–19.
- Ye, T.; Dong, L.; Xia, Y.; Sun, Y.; Zhu, Y.; Huang, G.; and Wei, F. 2025. Differential Transformer. In *International Conference on Learning Representations (ICLR)*.
- Yu, H.; Su, W.; Yang, Y.; liu, L.; Yuan, Y.; Xie, Y.; and Huang, T. 2025. A non-autoregressive Chinese-Braille translation approach with CTC loss optimization. *Expert Systems with Applications*, 269: 126356.
- Zhang, B.; and Sennrich, R. 2019. Root Mean Square Layer Normalization. In *Advances in neural information processing systems (NeurIPS)*, 12381 – 12392.