

Progress Ratio Embeddings: An Impatience Signal for Robust Length Control in Neural Text Generation

Ivanhoé Botcazou Tassadit Amghar Sylvain Lamprier Frédéric Saubion

LERIA, University of Angers

{ivanhoe.botcazou, tassadit.amghar, sylvain.lamprier, frederic.saubion}@univ-angers.fr

Abstract

Modern neural language models achieve high accuracy in text generation, yet precise control over generation length remains underdeveloped. In this paper, we first investigate a recent length control method based on *Reverse Positional Embeddings (RPE)* and show its limits when control is requested beyond the training distribution. In particular, using a discrete countdown signal tied to the absolute remaining token count leads to instability. To provide robust length control, we introduce *Progress Ratio Embeddings (PRE)*, as continuous embeddings tied to a trigonometric *impatience* signal. PRE integrates seamlessly into standard Transformer architectures, providing stable length fidelity without degrading text accuracy under standard evaluation metrics. We further show that PRE generalizes well to unseen target lengths. Experiments on two widely used news-summarization benchmarks and a popular question generation dataset validate these findings.¹

1 Introduction

Over the past few years, Language Models (LMs) have achieved remarkable progress in natural language processing, reaching state-of-the-art performance across a wide range of text generation tasks. Despite these advances, generating outputs that meet specific structural or stylistic constraints remains challenging. Among these constraints, length control is particularly fundamental, as the appropriate degree of compression or elaboration often depends on the task and the desired granularity of the information. In tasks such as summarization, for instance, the compression ratio between the source and the output naturally correlates with the number of words or tokens produced. Yet, controlling this aspect precisely remains both a technical and a cognitive challenge for neural language models (Kikuchi et al., 2016).

During inference, major LMs generate tokens sequentially until the process terminates by the sampling of the $\langle \text{EOS} \rangle$ token. However, this stopping criterion is inherently stochastic, and the decision mechanism guiding token selection at each step depends primarily on the previously generated context. As a result, the overall planning of the generation process lacks transparency, offering only limited opportunities for explicit control. In practice, users can request outputs of varying lengths, ranging from concise phrases to long, elaborated texts. For example, in abstractive summarization, one can look for either very short TL;DR-style summaries or slightly compressed paragraphs that preserve much of the original content. While instruction-following Large Language Models (LLMs) can express length preferences through prompting (Lou et al., 2024; Dong et al., 2024), this form of in-context control remains restricted. It is effective only for very large models trained on extensive amounts of general-purpose data. In this paper, we focus on the setting of Small Language Models (SMLs) trained on restricted, domain-specific datasets, where length control must instead be achieved through architectural design.

Recent approaches have addressed the problem of target length control in text generation, through the use of length-informed positional embeddings integrated into the inputs fed to the decoder of the LM. For example, (Miculicich et al., 2023; Butcher et al., 2025) propose to feed the Transformer decoder with reverse positional encodings that depend on the remaining number of tokens to be decoded given a desired target length. Although this enables alright length control for summaries whose lengths fall within a restricted window centered on the mode of the training distribution, in this paper, we show that it performs poorly for out-of-distribution target lengths.

To avoid this limitation, we introduce continuous *Progress Ratio Embeddings (PRE)*, which

¹Code available at: [Ivanbtz9/ProgressRatioEmbeddings](https://github.com/Ivanbtz9/ProgressRatioEmbeddings)

reduce the dependence on the discrete target length provided to the model. In contrast to discrete Reverse Positional Embeddings (RPE), PRE offers a smooth and continuous representation of progress that generalizes beyond the training distribution of reference lengths. PRE embeddings are introduced through trigonometric impatience signals whose frequency gradually increases until the end of generation. This design aligns the decoding process more closely with the user-specified target length and, by construction, makes it less dependent on the distribution of training reference lengths. We further integrate PRE into two notable pre-trained Encoder-Decoder architectures (Vaswani et al., 2017), achieving robust length invariance without compromising relevance or informativeness in neural text generation tasks.

Our main contributions are threefold:

- (1) We propose a novel length control strategy, called Progress Ratio Embeddings (PRE), based on continuous “impatience” signals that modulate the decoding process according to the completion progress;
- (2) We reproduce and adapt the Reverse Positional Embeddings (RPE) method on a classical encoder-decoder architecture, and provide an analysis of its strengths and weaknesses;
- (3) We demonstrate, on standard text summarization, that our PRE model can generalize length control beyond the training distribution, producing outputs of previously unseen lengths.

The paper is organized as follows. Section 2 presents related works. In Section 3 we introduce our method, called PRE. Section 4 describes the setup that is used for the experiments presented in Section 5 and Section 6.

2 Related Work

The large number of parameters and the probabilistic nature of LMs associate them with Black Box algorithms, which suffer from well-known weaknesses such as a lack of explicability and limited controllability in planning (Zhao et al., 2023). Given an input context, the vast space of possible continuations forces an autoregressive model to select a trajectory via sampling to maximize a context-conditioned probability.

To improve controllability, recent research has shown that LM outputs can be empirically enhanced through *chain-of-thought prompting* (CTP).

Within this perspective, existing study (Jie et al., 2024) propose prompt-based methods combined with reinforcement learning to train prompt extractors that enhance both performance and controllability. Moreover, by integrating prompting with iterative trajectory sampling under a Metropolis-Hastings framework, Gu et al. (2024) demonstrate how LLMs can be regulated to enforce constraints such as output length, without requiring parameter modifications. In this paper, we instead focus on smaller language models trained on domain-specific or resource-constrained datasets, where reference contents are distributed within a narrow length range. In such constrained settings, where contextual information about length is limited and prompting is not feasible, explicit architectural mechanisms become necessary to guarantee effective and reliable length control.

An alternative line of research focuses on modifying the decoding process itself. Decoding-based approaches, such as *LenAtten* (Yu et al., 2021), extend the autoregressive generation mechanism by conditioning the prediction head not only on the decoder’s hidden state but also on an additional length-context vector, thereby enabling finer-grained control over the generated output length. Beyond prompting and decoding, other studies introduce structural modifications to the model architecture. For example, the Length-Aware Attention Mechanism (LAAM), (Liu et al., 2022) operates at the attention level by gradually adjusting the cross-attention weights in a length-dependent manner, particularly between the <EOS> token in the input text and each decoder token (Appendix E).

In parallel, the remaining length is encoded via specialized positional representations, following the control mechanism introduced by Takase and Okazaki (2019) (Appendix G). Based on this concept, Miculicich et al. (2023) and Butcher et al. (2025) fine-tune pre-trained Transformers to incorporate RPE embeddings. At each decoding step, a discrete countdown index representing the remaining length is converted into a vector of the same dimension as the token and positional embeddings, and injected into the decoder (Appendix B). By training through next-token prediction, this additional signal enables the model to predict more accurately the expected stopping point.

However, the RPE method is explicitly dependent on the target length observed during training. As shown in Section 5, this dependency constitutes a critical limitation, RPE achieves correct control

only within the training length distribution, but globally fails to generalize for out-of-distribution lengths. This weakness highlights the need for novel architectural mechanisms that enable deterministic and more invariant length control in text generation. To address this challenge, we introduce the PRE method.

3 Continuous Length Control

In this section, we provide a brief overview of the role of positional embeddings in Transformer models. We then introduce our PRE method, which injects a lightweight *impatience signal* into the decoder inputs of a vanilla Transformer. We justify the theoretical foundation from the digital signal processing and expose the global fine-tuning aspects.

3.1 Transformers needs Positional Information

Transformer-based LMs process sequences of ordered tokens, where each token t is mapped to a token embedding² $E_t \in \mathbb{R}^{d_{\text{model}}}$. Because of the self-attention setup by itself, the input tokens are permutation invariant, the model must be given positional information. Each token index position can be mapped to a positional embedding $P_t \in \mathbb{R}^{d_{\text{model}}}$, which is added to the token embedding E_t before the first Transformer layer :

$$X_t = E_t + P_t. \quad (1)$$

This addition places content and position information in the same vector X_t and allows attention to condition on where a token occurs, not only on what it is. Depending on the model configuration, positional information can be encoded with fixed sinusoidal embeddings as initially introduced in (Vaswani et al., 2017). Alternatively, positional embeddings can be learned during the pre-training as presented in (Gehring et al., 2017; Radford et al., 2018). Some models incorporate relative positional embeddings that capture relative distance between sequence elements (Shaw et al., 2018). Modern LLMs typically use Rotary Positional Embeddings (RoPE), which are relative positional encodings applied inside each attention block (Jianlin et al., 2023). In all cases, without positional information, any permutation of the input tokens yields the same multiset of attention scores.

²An embedding is a dense vector of dimension d_{model}

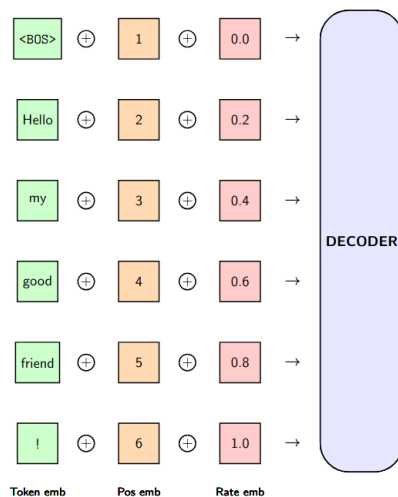


Figure 1: Illustration of Progress Ratio Embeddings added to token and positional embeddings to form the input of the decoder.

At the decoding step, the RPE methods additionally provide new position embeddings in reverse order to indicate how many tokens remain to be generated, rather than only giving the absolute position from the start of generation. When the model is trained with such reversed embeddings counting down from the initial target length to zero over the training reference the decoder block can modulate generation length at inference time according to a user-specified target length. However, we argue that this approach, while effective for lengths seen during training, does not generalize well to unseen target lengths, which are often required in practical applications. In the following, we propose a new method to overcome this limitation.

3.2 Progress Ratio Embeddings as an Impatience Signal

We introduce the PRE method based on progress ratios $r \in [0, 1]$ that quantify the advancement of the decoding process. At decoding step t for a requested target length l , we can define a progress ratio: $r_t = \frac{t}{l}$. As illustrated in Figure 1, a ratio-dependent embedding is added to the initial decoder input embeddings :

$$X_t = E_t + P_t + \xi(r_t) \quad (2)$$

where $\xi : [0, 1] \rightarrow \mathbb{R}^{d_{\text{model}}}$ denotes what we call the *Progress Ratio Embedding function*. A crucial advantage over the finite set of reversed positional embeddings used in RPE process is that PRE defines embeddings continuously for any ratio within the interval $[0, 1]$. This allows the model to compute

length-control information on the fly, covering the entire progress range rather than being restricted to discrete positions seen during training.

Our goal is to construct an embedding $\xi(r_t)$ that is rich enough to avoid overfitting to the specific ratios seen during training, that may differ at inference time when new target lengths are requested. We therefore seek a smooth representation that forms a meaningful metric space over the range $[0; 1]$. Specifically, we assume that for any possible ratio $r \in [0, 1]$, there exists a regular periodic function that represents the *impatience signal* of the process. By analogy, this signal can be seen as the impatience of a listener awaiting the end of a message. Building on signal sampling theory, we design a stress signal with pulsation ω_r (i.e., corresponding to a frequency $F = \omega_r/(2\pi)$) that increases with the progress ratio r , which enables the decoder to receive embeddings that encode the generation progress across $[0, 1]$. We consider in the following $\omega_r = r \times M$, with M an hyperparameter that scales the frequency increase of the signal (Appendix A).

The Nyquist–Shannon theorem (Shannon, 1949) guarantees that any periodic, band-limited function can be perfectly reconstructed from samples taken at a rate F_s exceeding twice its highest frequency component F_{max} . Furthermore, sinusoids provide a natural basis for such signals, the use of sine and cosine components ensure continuous phase encoding. Following this principle, we design our embedding so that each *pair* of dimensions encodes a single frequency component via cosine and sine in alternating dimensions. Formally, for $1 \leq j \leq d_{model}$,

$$\xi(r)_j = \begin{cases} \sin(\omega_r \cdot x_j), & \text{if } j \text{ is even,} \\ \cos(\omega_r \cdot x_j), & \text{if } j \text{ is odd,} \end{cases} \quad (3)$$

where ω_r controls the frequency of oscillations given by the progress ratio r . This formulation uses $x_j = \frac{2 \lfloor j/2 \rfloor}{d_{model}}$ as the independent variable for each dimension j , which corresponds to a sampling rate of $F_s = d_{model}/2$ given points regularly spaced over $[0, 1]$. Taking $M = d_{model} \cdot \frac{\pi}{2}$ allows to benefit from the full capacity of the considered representation, while ensuring that the Nyquist–Shannon condition is satisfied, since we have in that case: $F = \frac{\omega_r}{2\pi} \leq 1 \cdot d_{model}/4 = \frac{d_{model}}{4}$, and thus: $F_s = \frac{d_{model}}{2} = 2F_{max}$. In practice, we use $M = \frac{d_{model}}{2}$, leaving a small safety margin below the Nyquist bound to accommodate rounding

and numerical precision limits.

3.3 Finetuning with PRE

Let $A := (A_0, \dots, A_n)$ denote the sequence of tokens corresponding to the context of the task (e.g., the source content), and $S := (S_0, \dots, S_l)$ the sequence of tokens in a "gold" reference (e.g., a target text), where l is the expected number of tokens in the produced text. Following Section 3.2, we consider an associated family of Progress Ratio Embeddings $\Xi = (\xi(r_0), \dots, \xi(r_l))$ related to the expected target length l . Given a pretrained model³, our objective is to adapt its parameters so that it predicts the conditional probability $P(S | A, l)$.

Following the maximum likelihood estimation (MLE) principle, we optimize the model parameters θ to maximize the product of conditional probabilities. This product can be factorized using the chain rule (\star) and the causal property ($\star\star$) of the decoder:

$$\begin{aligned} P_\theta(S | A, \Xi) &\stackrel{\star}{=} \prod_{t=0}^l P_\theta(S_t | S_{<t}, A, \Xi) \\ &\stackrel{\star\star}{=} \prod_{t=0}^l P_\theta(S_t | S_{<t}, A, \Xi_{<t}). \end{aligned}$$

The training objective is defined by the standard cross-entropy loss, applied over mini-batches of size m ⁴. We minimize the empirical loss over a batch $\mathcal{B} = \{(A^i, S^i, \Xi^i)\}_{1 \leq i \leq m}$, defined as the average negative log-likelihood across all tokens.

$$\mathcal{L}_{\mathcal{B}}(\theta) = -\frac{1}{m} \sum_{i=1}^m \sum_{t=0}^{l_i} \log P_\theta(S_t^i | S_{<t}^i, A^i, \Xi_{<t}^i)$$

Minimizing the sequence loss while conditioning the decoder on a normalized progress ratio provides a consistent source of temporal side information. To enhance generalization, Gaussian noise is injected into each ratio (Eq. 4), exposing the model to a continuous spectrum of values over $[0, 1]$ and promoting smooth interpolation rather than dependence on a few discrete targets.

$$r = \text{Clip} \left(r + \frac{2\delta}{d_{model}} ; 0 ; 1 \right) \quad (4)$$

where $\delta \sim \mathcal{N}(0, 1)$

³In our experiments, we focus on encoder-decoder models, although our technique can be straightforwardly applied to decoder-only architectures by initializing the progress ratios at the beginning of the expected sequence S .

⁴Computational resources allowed us to use $m = 5$

Once the generated sequence reaches or exceeds the target length, the ratio is clipped at 1 causing the impatience signal to saturate and thereby discouraging further generation.

4 Experimental Setup

In this section, we present and justify our choices of model, datasets, hyperparameters, and evaluation metrics as integral components of the PRE method’s implementation and evaluation.

4.1 Model

For our experimentations, we used the proven BART encoder-decoder model as backbone (Lewis et al., 2019), it achieves strong performances on specific tasks and remains highly suitable for fine-tuning. As a sequence-to-sequence model pre-trained with a denoising objective inspired by BERT (Devlin et al., 2019), it has been shown to be effective in a range of NLP applications, including abstractive summarization and question generation task. Released by Meta AI with about 400M parameters⁵, it can be trained on standard computing resources and supports inputs of up to 1,024 tokens.

4.2 Datasets

We conduct experiments on two widely used news-summarization benchmarks and on a question-answering dataset.

CNN/DM (Nallapati et al., 2016) contains articles paired with summaries that are typically closely aligned with source sentences.

XSum (Narayan et al., 2018) provides highly abstractive one-sentence summaries that require substantial rephrasing beyond the source text.

SQuAD (Rajpurkar et al., 2016) is the Stanford Question Answering Dataset, consisting of questions posed on Wikipedia articles.

4.3 Metrics

For each input article, there may exist multiple human-judged acceptable candidates that preserve the same level of information. Due to the existence of these valid alternatives, identifying metrics that accurately capture the relevance and informativeness of a generated summary remains highly challenging (Koh et al., 2022).

Given an input text and a gold-standard output summary, the most common metrics used to eval-

uate the relevance of abstractive candidates are ROUGE-1, ROUGE-2, and ROUGE-L (Lin, 2004), which provide an F1-score based on n -gram overlap between the generated output and the reference summary. However, these metrics are often questioned, as they do not correlate strongly with human judgments (Deutsch et al., 2022). Despite this limitation, ROUGE remains the primary score reported in a large number of studies, including ours. Acknowledging the weaknesses of ROUGE, the BERTScore (B.S.) (Zhang et al., 2020) has been proposed to avoid this lack of relevance. To further extend our evaluation, we employ a fine-grained LLM-as-a-judge framework specifically tailored for summarization. This approach is inspired by the FineSurE evaluation methodology (Song et al., 2024), with comprehensive implementation details provided in (Appendix H). We evaluate length controllability using the mean absolute error (MAE) between the generated and target lengths. Unless otherwise stated, lengths are measured in tokens as defined by each model’s tokenizer.

4.4 Hyperparameters and Devices

All models were trained on the full training set using the AdamW optimizer (Loshchilov and Hutter, 2019), with hyperparameters $\beta_1 = 0.9$, $\beta_2 = 0.99$, and $\epsilon = 1 \times 10^{-8}$. The learning rate was set to 1×10^{-5} . All reported experiments and results are evaluated on the full test split of each dataset.

5 RPE results and limitations of the summary task

We begin by reproducing the RPE method with BART-L and fine-tuning it on two famous summary benchmark datasets. This reproduction confirms the reported length-control behavior, as shown in Table 1, consistently regarding results from the original paper (Miculicich et al., 2023). It also indicates this method maintains strong summarization quality relative to baselines without length control, which we reported in Table 2. RPE-BART-L achieves accurate length control within each 10-token target-length bucket.

However, in approximately 2% of cases, the MAE between the generated and expected lengths exceeds 10 tokens (see outliers in Figure 2). Moreover, as shown in Figure 3, when the requested target length exceeds 350 tokens, the effectiveness of length control degrades markedly.

⁵We use the BART large version, where $d_{\text{model}} = 1,024$

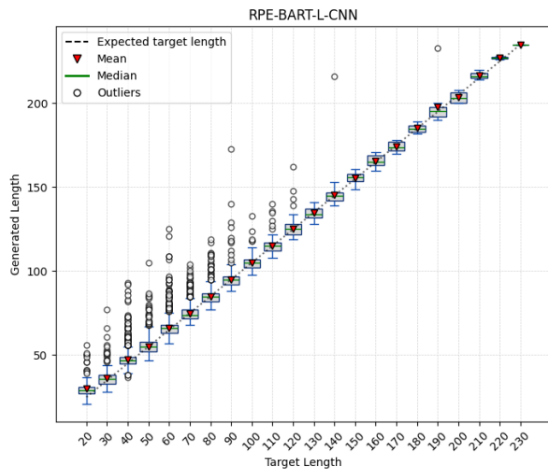


Figure 2: MAE by target-length bucket (10 tokens) for RPE-BART-L on CNN/DailyMail.

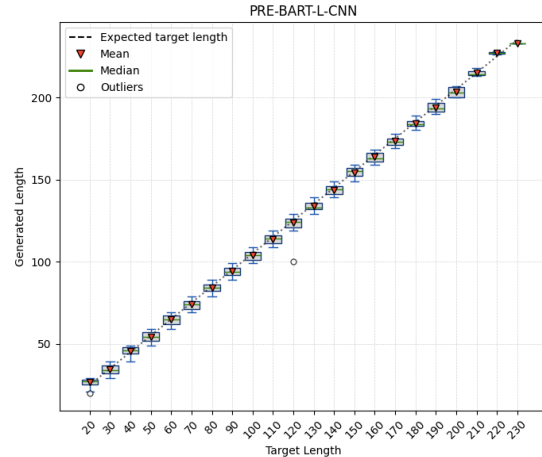


Figure 4: MAE by target-length bucket (10 tokens) for PRE-BART-L on CNN/DailyMail.




Figure 3: MAE by target-length bucket (25 tokens) for RPE-BART-L on CNN/DailyMail when target lengths above 300 tokens are requested.

This lack of generalization⁶ is precisely what motivated the introduction of our - length invariant - PRE approach, that we experiment in the next section.

6 PRE results and analysis

Summaries under length control

Results reported in Table 2 demonstrate that our PRE method consistently yields high-quality summaries while maintaining precise length control. Specifically, the PRE approach sustains high ROUGE and BERTScore performance by achieving a low Mean Absolute Error (MAE) between target and generated lengths across the full spectrum

⁶Target lengths chosen according to the input article length and sampled at random from out-of-distribution lengths.

of length constraints (see Figures 4 and 14). Table 1 further contrasts PRE with both no-control baselines and our RPE reimplementations, confirming that PRE delivers superior controllability without sacrificing quality.

Dataset	Model	Type	MAE [↓] ± SD
CNN/DM	BART-L	N	19.2 ± 17
	RPE-BART-L	G	1.6 ± 3.6
	PRE-BART-L	G	0.5 ± 0.3
XSum	BART-L	N	5.8 ± 5
	RPE-BART-L	G	0.7 ± 1.1
	PRE-BART-L	G	0.1 ± 0.2

Table 1: Length control evaluation on CNN/DailyMail and XSum using MAE and standard deviation (SD). **Legend:** G = reference length; N = no explicit length control. Lower values indicate better control.

CNN/DailyMail				
Model	Type	R-1 [↑]	R-2 [↑]	R-L [↑] B.S. [↑]
BART-L (Lewis et al., 2019)	N	44.2	21.1	40.9 69.7
RPE-BART-L	G	44.5	21.2	41.3 69.4
PRE-BART-L	G	45.3	21.9	42.2 69.8
XSum				
Model	Type	R-1 [↑]	R-2 [↑]	R-L [↑] B.S. [↑]
BART-L (Lewis et al., 2019)	N	45.14	22.27	37.25 73.3
RPE-BART-L	G	44.5	20.8	35.6 72.3
PRE-BART-L	G	45.2	21.3	36.4 72.7

Table 2: ROUGE and BERTScore results on CNN/DailyMail and XSum. **Legend:** G = length from the reference summary is given; N = no explicit length control.

Generalization to out-of-distribution Data

We have introduced the PRE method to achieve more robust generalization of length control across the full range of lengths that the model can handle. Figure 5 illustrates that PRE-BART-L is actually able to generate summaries with accurate length control, even when the requested output length exceeds 300 tokens.




Figure 5: MAE by target-length bucket (25 tokens) for PRE-BART-L on CNN/DailyMail when target lengths above 300 tokens are requested.

To evaluate behavior along all manageable context windows, we also investigate a comparative analysis between PRE and RPE methods on 6,000 randomly selected examples from the CNN/Daily Mail test set. For each selected article, we randomly assigned a target summary length smaller than the article length and asked both models to generate a summary under this constraint. Table 3 reports corresponding results, in terms of the percentage of produced outliers for which MAE greater than 20 tokens (% Out). We observe that the proportion of outliers is considerably smaller with PRE, highlighting its robustness in handling long and variable-length generation tasks. PRE encodes a continuous ratio signal normalized between 0 and 1. This distinction is crucial for handling arbitrary user-specified output lengths. RPE is more sensitive to the distribution of target lengths in the training data. As shown in Appendix C, long target lengths are underrepresented in the original dataset. Consequently, models trained with RPE rarely encounter such cases during fine-tuning, thus their parameters are not adapted to accommodate very long summaries.

Bucket	% Out RPE ↓	% Out PRE ↓	Count
[300, 350)	10.8%	0.4%	839 / 922
[350, 400)	15.4%	1.0%	705 / 809
[400, 450)	22.3%	1.5%	608 / 745
[450, 500)	28.0%	3.4%	453 / 629
[500, 550)	37.0%	4.3%	357 / 577
[550, 600)	50.1%	5.4%	245 / 490
[600, 650)	51.4%	4.0%	216 / 429
[650, 700)	53.9%	7.4%	179 / 374
[700, 750)	64.8%	10.9%	124 / 351
[750, 800)	70.3%	9.9%	93 / 273
[800, 850)	75.1%	9.0%	76 / 242
[850, 900)	81.3%	8.4%	43 / 230
[900, 950)	84.9%	8.9%	35 / 214
[950, 1000)	95.8%	19.8%	4 / 65

Table 3: Comparison of outlier rate between RPE-BART-L and PRE-BART-L across length buckets. The Count column reports, respectively, the number of non-outlier samples for each model.

Question generation accomodation

To demonstrate the strong generalization capability of our method across diverse neural text generation tasks, we conducted experiments on the SQuAD dataset. We fine-tune the BART model with both RPE and PRE for the question generation task, where the objective is to generate a relevant question given a context passage and its associated answer. As shown in Tables 4 and 5, the PRE variant exhibits clear advantages: it enables highly accurate length-controlled question generation while preserving the semantic quality of the outputs. Moreover, our method achieves strong results across multiple evaluation metrics, including BLEU, ROUGE, and BERTScore.

Model	Type	MAE [↓] ± SD
BART-L	N	3.12 ± 3.3
RPE-BART-L	G	0.8 ± 3.6
PRE-BART-L	G	0.0 ± 0.1

Table 4: Length control evaluation on SQuAD using MAE and standard deviation (SD). **Legend:** G = reference length; N = no explicit length control. Lower values indicate better control.

These experiments demonstrate the scalability of the PRE method to a broad range of text generation tasks. In particular, PRE provides more precise control over the target length than the RPE method (Figures 6 and 7), especially near the extremities of the desired length distribution observed during training. This highlights the robustness and reli-

ability of PRE for fine-grained length-controlled generation, even under demanding representation constraints.

SQuAD			
Score	BART-L	RPE-BART-L	PRE-BART-L
BLEU [↑]	16.7	16.4	18.6
R-1 [↑]	54.4	52.1	55.3
R-2 [↑]	32	30.1	32.9
R-L [↑]	50.1	47.7	50.8
B.S. [↑]	75.2	72.6	74.8

Table 5: BLEU, ROUGE and BERTScore results on the SQuAD dataset.




Figure 6: MAE by target-length bucket (1 tokens) for RPE-BART-L on SQuAD dataset.




Figure 7: MAE by target-length bucket (2 tokens) for PRE-BART-L on SQuAD dataset.

7 Conclusion

In this paper, we have introduced Progress Ratio Embeddings (PRE) as a novel and robust mecha-

nism for length control in neural text generation. Unlike traditional Reverse Positional Embeddings (RPE), which rely on discrete countdown signals and struggle with generalization beyond the training distribution, PRE uses continuous impatience signals to encode progress in a more scalable way. Our experiments across different tasks generation highlight that PRE significantly improves length controllability while maintaining or enhancing text generation quality, according to classic metrics. Furthermore, PRE enables generalization to out-of-distribution target lengths, reducing the proportion of outliers and mean absolute error compared to RPE. Beyond providing length control in abstractive summarization (Almohaimeed and Azmi, 2025), PRE opens up new directions for fine-grained control over the generation process, enabling more progressive summarization while improving explainability.

8 Limitations

Several limitations can be addressed in our work. Let us briefly discuss them. In this article, we have only describe our PRE method within an encoder-decoder architecture, which is not the prevailing trend of current decoder-only language models. A natural next step would be to evaluate the robustness of length control with the PRE method on decoder-only models and explore how this length control could facilitate step-by-step reasoning in content production, especially when given a progressively decreasing target length. Moreover, length control is not only relevant for summarization and question generation tasks but many other text-to-text tasks could benefit from this capability to regulate the generated length. Finally, an interesting direction for future research would be to explore length control within *chain-of-thought* models. Being able to regulate the reasoning length could make such models more efficient at inference time, by preventing unnecessary over-generation and by allowing a deterministic adjustment of the reasoning depth.

9 Ethical Considerations

We have designed our methods to be simple to reproduce as they rely on publicly available pre-trained language models. We acknowledge that any text generation system carries potential risks of producing biased or misleading content. To mitigate these risks, we restricted our experiments to widely

used, publicly datasets whose content is considered as relatively safe for research purposes. Our work complies with the ACL ETHICS POLICY ⁷.

10 LLM Usage Declaration

During the preparation of this manuscript, we used Large Language Models to assist with text clarity, grammar, and formulation, particularly in polishing the abstract and certain explanatory sentences. The scientific content, experimental design, results, and interpretations were entirely conceived and written by the authors. LLMs were not used to generate original ideas, proofs, or analyses; its contribution was limited to language refinement.

11 Acknowledgements

This work benefited from access to the HPC resources of GLICID ⁸ and resources of IDRIS ⁹ under the allocation 2025-AD011016042 and 2026-AD011014032R3 made by GENCI.

References

- Norah Almohaimed and Aqil M. Azmi. 2025. *Abstractive text summarization: A comprehensive survey of techniques, systems, and challenges*. *Computer Science Review*, 57:100762.
- Bradley Butcher, Michael O’Keefe, and James Titchener. 2025. *Precise length control for large language models*. *Natural Language Processing Journal*, 11:100143.
- Daniel Deutsch, Rotem Dror, and Dan Roth. 2022. *Re-Examining System-Level Correlations of Automatic Summarization Evaluation Metrics*. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 6038–6052. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. *Bert: Pre-training of deep bidirectional transformers for language understanding*. *Preprint*, arXiv:1810.04805.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Baobao Chang, Xu Sun, Lei Li, and Zhifang Sui. 2024. *A survey on in-context learning*. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1107–1128, Miami, Florida, USA. Association for Computational Linguistics.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. *Convolutional sequence to sequence learning*. *Preprint*, arXiv:1705.03122.
- Google. Google. <https://research.google/>. Accessed: 2025-09-22.
- Yuxuan Gu, Wenjie Wang, Xiaocheng Feng, Weihong Zhong, Kun Zhu, Lei Huang, Tat-Seng Chua, and Bing Qin. 2024. *Length controlled generation for black-box llms*. *Preprint*, arXiv:2412.14656.
- Renlong Jie, Xiaojun Meng, Lifeng Shang, Xin Jiang, and Qun Liu. 2024. *Prompt-Based Length Controlled Generation with Multiple Control Types*. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 1067–1085. Association for Computational Linguistics.
- Yuta Kikuchi, Graham Neubig, Ryohei Sasano, Hiroya Takamura, and Manabu Okumura. 2016. *Controlling Output Length in Neural Encoder-Decoders*. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1328–1338. Association for Computational Linguistics.
- Huan Yee Koh, Jiabin Ju, Ming Liu, and Shirui Pan. 2022. *An empirical survey on long document summarization: Datasets, models, and metrics*. *ACM Computing Surveys*, 55(8):1–35.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. *Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension*. *Preprint*, arXiv:1910.13461.
- Chin-Yew Lin. 2004. *ROUGE: A package for automatic evaluation of summaries*. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Yizhu Liu, Qi Jia, and Kenny Zhu. 2022. *LAAM : Length Control in Abstractive Summarization by Pre-training Information Selection*. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6885–6895. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2019. *Decoupled weight decay regularization*. *Preprint*, arXiv:1711.05101.
- Renze Lou, Kai Zhang, and Wenpeng Yin. 2024. *Large language model instruction following: A survey of progresses and challenges*. *Computational Linguistics*, 50(3):1053–1095.
- Meta AI. Meta ai research. <https://ai.meta.com/>. Accessed: 2025-09-22.

⁷<https://www.aclweb.org/portal/content/acl-code-ethics>

⁸<https://glicid.fr/>

⁹<https://www.genci.fr/>

Lesly Miculicich, Yujia Xie, Song Wang, and Pengcheng He. 2023. [REPILOT : Summarization with Precise Length Control](#). *Preprint*, arXiv:2305.05171.

Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gulçehre, and Bing Xiang. 2016. [Abstractive text summarization using sequence-to-sequence RNNs and beyond](#). In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290, Berlin, Germany. Association for Computational Linguistics.

Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. [Don't give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807, Brussels, Belgium. Association for Computational Linguistics.

Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. [Improving Language Understanding by Generative Pre-Training](#).

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2023. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Preprint*, arXiv:1910.10683.

Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [Squad: 100,000+ questions for machine comprehension of text](#). *Preprint*, arXiv:1606.05250.

C.E. Shannon. 1949. [Communication in the presence of noise](#). *Proceedings of the IRE*, 37(1):10–21.

Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. [Self-attention with relative position representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 464–468, New Orleans, Louisiana. Association for Computational Linguistics.

Shlomo Stept. 2023. [T5-large fine-tuned on xsum + cnn/dailymail for abstractive summarization](#).

Sho Takase and Naoaki Okazaki. 2019. [Positional Encoding to Control Output Sequence Length](#). *Preprint*, arXiv:1904.07418.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

Zhongyi Yu, Zhenghao Wu, Hao Zheng, Zhe XuanYuan, Jefferson Fong, and Weifeng Su. 2021. [LenAtten: An Effective Length Controlling Unit For Text Summarization](#). *Preprint*, arXiv:2106.00316.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [BERTScore: Evaluating Text Generation with BERT](#). *Preprint*, arXiv:1904.09675.

Haiyan Zhao, Hanjie Chen, Fan Yang, Ninghao Liu, Huiqi Deng, Hengyi Cai, Shuaiqi Wang, Dawei Yin, and Mengnan Du. 2023. [Explainability for large language models: A survey](#). *Preprint*, arXiv:2309.01029.

Jianlin Su and Yu Lu and Shengfeng Pan and Ahmed Murtadha and Bo Wen and Yunfeng Liu. 2023. [RoFormer: Enhanced Transformer with Rotary Position Embedding](#). *Preprint*, arXiv:2104.09864.

Hwanjun Song, Hang Su, Igor Shalyminov, Jason Cai, and Saab Mansour. 2024. [FineSurE: Fine-grained Summarization Evaluation using LLMs](#). *Preprint*, arXiv:2407.00908.

Gemma Team. 2025. [Gemma 3. Technical Report](#), Kaggle.

A Appendix : Stress signal plot for varying pulsations

We use trigonometric functions to represent the concept of a periodic impatience signal as:

$$s : \mathbb{R}^+ \times [0, 1] \longrightarrow [-1, 1]^2$$

$$(\omega, x) \longmapsto (\cos(\omega x), \sin(\omega x)),$$

where ω denotes the pulsation of the periodic signal with a corresponding frequency $F = \frac{\omega}{2\pi}$.




Figure 8: Impatience signal curves plot for different pulsation ω based on $M = \frac{BART_{d_{model}}}{2} = 512$.

B Appendix : Illustration of the RPE method in relation to the decoder block input

Reverse Positional Embeddings (RPE) extend the standard sinusoidal formulation (Vaswani et al., 2017) by explicitly encoding the number of tokens remaining until the desired target length. Instead of encoding the absolute position i , RPE use the countdown index $(l-i)$, where l is the target length.

$$RPE(i, 2k) = \sin\left(\frac{l-i}{10000^{2k/d_{model}}}\right), \quad (5)$$

$$RPE(i, 2k+1) = \cos\left(\frac{l-i}{10000^{2k/d_{model}}}\right), \quad (6)$$

where d_{model} is the embedding dimension and k indexes the dimensions. This design injects a discrete countdown signal into the decoder, enabling the model to condition generation explicitly on the remaining length.

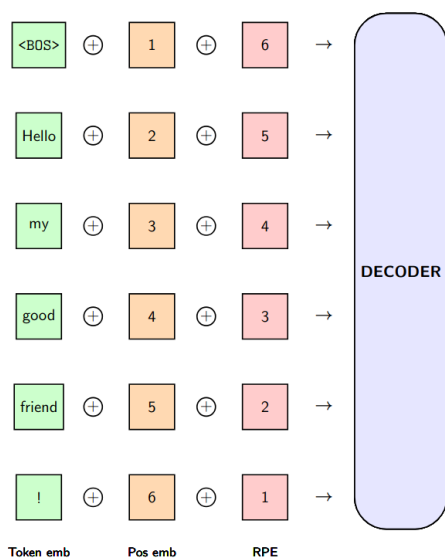


Figure 9: Illustration of Reverse Positional Embeddings (RPE) added to token and positional embeddings before being passed to the decoder block.

C Appendix : Summary lengths distribution through datasets

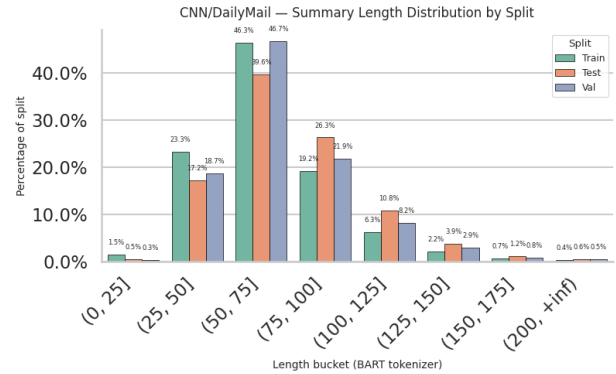


Figure 10: Data visualisation over the CNN/DailyMail summary length distribution.




Figure 11: Data visualisation over the XSum summary length distribution.

D Appendix : Results of the PRE method adaptation on the T5 encoder-decoder architecture

The T5-Large model (Raffel et al., 2023) treats every NLP task as text-to-text. Pre-trained on the C4 corpus with a span-corruption objective, it contains roughly 770M parameters and accepts inputs of up to 512 tokens. Developed by Google, this open-source model offers strong transferability across tasks and remains competitive for abstractive summarization. For this model $d_{model} = 512$.

Dataset	Model	Type	MAE [↓] ± SD
CNN/DM	T5-L	N	20.1 ± 17.9
	PRE-T5-L	G	7.9 ± 7.8
XSum	T5-L	N	6.5 ± 5.4
	PRE-T5-L	G	1.9 ± 3.3

Table 6: Length control evaluation on CNN/DailyMail and XSum using MAE and standard deviation (SD). **Legend:** G = guided by reference length; N = no explicit length control. Lower values indicate better control.

CNN/DailyMail					
Model	Type	R-1 [↑]	R-2 [↑]	R-L [↑]	B.S. [↑]
T5-L (Raffel et al., 2023)	N	42.50	20.68	39.75	68.2
PRE-T5-L	G	44.6	21.0	36.0	69.4
XSum					
Model	Type	R-1 [↑]	R-2 [↑]	R-L [↑]	B.S. [↑]
T5-L (Stept, 2023)	N	36.77	14.69	30.06	68.9
PRE-T5-L	G	42.8	19.1	34.3	71.7

Table 7: ROUGE and BERTScore results on CNN/DailyMail and XSum for T5 models. **Legend:** G = length from the reference summary is given; N = no explicit length control.

E Appendix: Reproduction of the LAAM Method

The Length-Aware Attention Mechanism (LAAM) (Liu et al., 2022) modifies the cross-attention of a seq2seq model by boosting the top- ℓ_t attended encoder tokens according to the remaining target length. This encourages the encoder-decoder attention to select information compatible with the desired output size. In our reproduction, LAAM is applied only to the last decoder layer of BART and fine-tuned on CNN/DailyMail and on SQuAD. Below we report the full set of scores obtained from our implementation.

Stat	R-1	R-2	R-L	B.S.	MAE
Mean	44.2	21.1	41.1	69.6	17.5
Std	12.5	13.6	12.6	6.9	16.2
Min	0.0	0.0	0.0	42.4	0.0
25%	36.0	11.6	32.7	65.4	5.0
50%	44.0	18.8	40.5	69.6	10.0
75%	51.9	27.9	48.5	74.0	21.0

Table 8: LAAM performance on CNN/DailyMail.

Stat	R-1	R-2	R-L	BLEU	B.S.	MAE
Mean	54.7	32.1	50.2	16.8	75.4	3.12
Std	23.4	27.6	24.3	27.1	12.4	3.18
Min	0.0	0.0	0.0	0.0	38.9	0.0
25%	37.5	10.0	31.3	0.0	66.5	1.0
50%	54.5	26.7	47.7	0.0	75.1	2.0
75%	71.4	50.0	66.7	30.2	84.2	4.0

Table 9: LAAM performance on SQuAD.




Figure 12: MAE by target-length bucket (25 tokens) for LAAM-BART-L on CNN/DailyMail dataset.

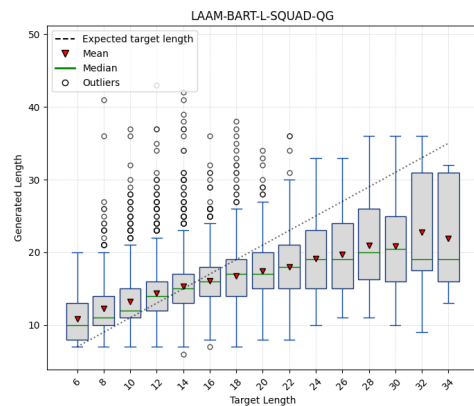


Figure 13: MAE by target-length bucket (2 tokens) for LAAM-BART-L on SQuAD dataset.

F Appendix: Statistical Significance of Length Fidelity

Statistical Significance (Paired t-tests). We performed paired Student’s t -tests on example-level MAE values from the CNN/DailyMail test set to assess whether the improvements of PRE over existing baselines are statistically reliable.

- PRE-BART vs. BART baseline:
 $t = -98.95, p < 10^{-300}$
- PRE-BART vs. RPE-BART:
 $t = -19.83, p = 4.8 \times 10^{-86}$

These results demonstrate that PRE yields highly statistically significant improvements in length fidelity.

G Appendix: Positioning with Length-Ratio Positional Encoding (LRPE)

Prior approaches such as the length-ratio positional encoding (LRPE) introduced by Takase and Okazaki (2019), which we initially cited as a source of inspiration proposes a formulation based on a ratio encoding during the decoding progress. We argue that, similarly to RPE, it still depends explicitly on the target length and we claim that this dependence hinders generalization, particularly for out-of-distribution target lengths at test time.

The LRPE formulation can be rewritten by introducing a progress ratio $r = \frac{pos}{len}$, similar to the one used in our approach:

$$\begin{aligned} \text{LRPE}(pos, len, 2j) &= \cos\left(\frac{pos}{len^{\frac{2j}{d_{\text{model}}}}}\right) \\ &= \cos\left(r \cdot len^{1-\frac{2j}{d_{\text{model}}}}\right), \\ \text{LRPE}(pos, len, 2j+1) &= \sin\left(\frac{pos}{len^{\frac{2j+1}{d_{\text{model}}}}}\right) \\ &= \sin\left(r \cdot len^{1-\frac{2j+1}{d_{\text{model}}}}\right). \end{aligned}$$

where j indexes the embedding dimension (for a model of size d_{model}), len denotes the target length, and pos is the current decoding step.

Although LRPE incorporates a progress ratio, it remains explicitly dependent on len . During training, this implies that the inputs are not invariant to the target length. In particular, the same ratio yields different encodings depending on whether the expected sequence is long or short.

To empirically validate this hypothesis, we reproduced the LRPE method and scaled it to the BART-L architecture on the CNN/DailyMail dataset. As shown in Tables 11 and 12, LRPE yields less stable results and does not reach the same level of accuracy as PRE in our main experiments.

To evaluate robustness under strict length constraints outside the training distribution, we compute the Mean Absolute Error (MAE) between the generated length and randomly sampled target lengths that exceed the original reference length. Results are grouped into 25-token buckets.¹⁰

As shown in Table 10, PRE-BART-L consistently achieves the lowest MAE across all length

¹⁰The same random seed is used across all models for fair comparison.

Bucket	PRE-BART-L	RPE-BART-L	LRPE-BART-L
[200, 225)	0.54	2.61	13.22
[225, 250)	0.63	3.18	16.91
[250, 275)	0.70	4.34	21.21
[275, 300)	0.87	8.60	26.72
[300, 325)	1.32	11.77	26.54
[325, 350)	1.19	10.71	30.95
[350, 375)	1.96	15.47	30.47
[375, 400)	1.58	20.69	34.04
[400, 425)	2.07	29.65	27.74
[425, 450)	3.20	37.02	26.51
[450, 475)	2.79	50.86	32.44
[475, 500)	2.96	57.29	43.14
[500, 525)	3.45	51.93	53.93
[525, 550)	4.03	58.24	58.28
[550, 575)	8.22	72.94	59.41
[575, 600)	7.91	94.44	46.09
[600, 625)	4.97	97.74	50.47
[625, 650)	6.83	113.09	45.10
[650, 675)	6.83	140.51	75.16
[675, 700)	12.78	143.27	86.23
[700, 725)	11.76	165.07	84.73
[725, 750)	8.48	170.40	117.84
[750, 775)	11.81	194.53	157.12
[775, 800)	7.33	164.31	141.47
[800, 825)	12.56	188.24	128.55
[825, 850)	7.60	242.80	126.77
[850, 875)	10.07	257.50	123.00
[875, 900)	16.88	266.51	137.03
[900, 925)	14.44	252.15	121.90
[925, 950)	9.49	329.41	105.84

Table 10: MAE by target-length bucket in out-of-distribution settings. Lower values indicate better length control.

ranges, including extreme out-of-distribution regimes.

In contrast, RPE-BART-L exhibits a sharp and near-monotonic degradation as the target length increases, highlighting its limited extrapolation capability. LRPE-BART-L improves over RPE but still shows substantial error growth for longer sequences, indicating residual dependence on absolute length information.

Overall, these results demonstrate that PRE provides significantly more stable and scalable length

Dataset	Model	Type	MAE \downarrow \pm SD
CNN/DM	BART-L	N	19.2 \pm 17
	LRPE-BART-L	G	2.19 \pm 1.89
	RPE-BART-L	G	1.6 \pm 3.6
	PRE-BART-L	G	0.5 \pm 0.3

Table 11: Length control evaluation on CNN/DailyMail using mean absolute error (MAE) and standard deviation (SD). **Legend:** G = reference length provided; N = no explicit length control. Lower values indicate better control.

control, making it the most reliable approach for generation under unseen length constraints.

CNN/DailyMail					
Model	Type	R-1 [↑]	R-2 [↑]	R-L [↑]	B.S. [↑]
BART-L	N	44.2	21.1	40.9	69.7
LRPE-BART-L	G	43.2	20.1	40.2	68.7
RPE-BART-L	G	44.5	21.2	41.3	69.4
PRE-BART-L	G	45.3	21.9	42.2	69.8

Table 12: ROUGE and BERTScore results on CNN/DailyMail.

H Appendix: Evaluation with LLM-as-a-Judge (FineSurE)

We conducted additional experiments using an LLM-as-a-judge evaluation metric, namely *FineSurE* (Fine-grained Summarization Evaluation) (Song et al., 2024). As the reference evaluator, we used GEMMA-3-12B-Instruct (Gemma Team, 2025)¹¹. Using the CNN/DailyMail test set, for each reference document–summary pair (A, S) , we computed the difference between the reference faithfulness score $\text{FineSurE}_{\text{Faithfulness}}(A, S)$ and the generated faithfulness score $\text{FineSurE}_{\text{Faithfulness}}(A, S^*)$, where S^* denotes the summary generated by each model (BART-L, RPE-BART-L, PRE-BART-L).

We then measured the absolute difference:

$$\Delta_{\text{Faith}} = |\text{Faithfulness}(A, S) - \text{Faithfulness}(A, S^*)|,$$

where lower values indicate that the generated summary is closer to the reference in terms of faithfulness. Our results show that PRE-BART-L produces faithfulness scores that are substantially closer to the reference compared to the two baseline models.


Dataset	Model	$\Delta_{\text{Faith}} \downarrow \pm \text{SD}$
CNN/DM	BART-L	0.12 ± 0.25
	RPE-BART-L	0.09 ± 0.26
	PRE-BART-L	0.01 ± 0.28

Table 13: Faithfulness distance to the reference using FineSurE (lower is better).


I Appendix : Density comparison of summary length distributions

The PRE method produces summaries whose length distribution closely matches the true distribution observed in the reference data.


¹¹<https://huggingface.co/google/gemma-3-12b-it>



(a) Reference length distribution.



(b) BART-L length distribution.



(c) PRE-BART-L length distribution.

Figure 14: Density comparison of summary length distributions on CNN/DailyMail.