

Evade the Trap of Mediocrity: Promoting Diversity and Novelty in Text Generation via Concentrating Attention

Wenhao Li^{1,2,3}, Xiaoyuan Yi⁵, Jinyi Hu^{1,2,3}, Maosong Sun^{1,2,3,4*}, Xing Xie⁵

¹ Department of Computer Science and Technology, Tsinghua University, Beijing

² Beijing National Research Center for Information Science and Technology

³ Institute for Artificial Intelligence, Tsinghua University, Beijing

⁴ Jiangsu Collaborative Innovation Center for Language Ability, Jiangsu Normal University, Xuzhou

⁵ Microsoft Research Asia

wh-li20@mails.tsinghua.edu.cn, xiaoyuanyi@microsoft.com

Abstract

Recently, powerful Transformer architectures have proven superior in generating high-quality sentences. Nevertheless, these models tend to produce dull high-frequency phrases, severely hurting the diversity and novelty of generated text. In this work, we dig into the intrinsic mechanism of this problem and found that sparser attention values in Transformer could improve diversity. To understand such a phenomenon, we first conduct both empirical and theoretical analysis and then attribute it to representation degeneration caused by the attentive mixture of the hidden states during training. We term this process the *Trap of Mediocrity*. To escape from such a trap, we introduce a novel attention regularization loss to control the sharpness of the attention distribution, which is transparent to model structures and can be easily implemented within 20 lines of python code. We prove that this method could be mathematically regarded as learning a Bayesian approximation of posterior attention. Experiments show that our method improved the diversity and novelty of the generated text while maintaining comparable quality on a variety of conditional and unconditional generation tasks.

1 Introduction

Natural Language Generation (NLG) plays a crucial role in modern natural language Processing (NLP). With considerable benefits for a wide array of applications such as automatic literary writing (Yi et al., 2018; Ao et al., 2021), summarization (Huang et al., 2021) and paraphrasing (Lin et al., 2021), this area has risen to prominence.

Recently, neural architectures based on the prevalent Transformer (Vaswani et al., 2017), e.g., GPT-2 (Radford et al., 2019), have proven highly successful in producing fluent text indistinguishable from the human-written one. Nevertheless, these models

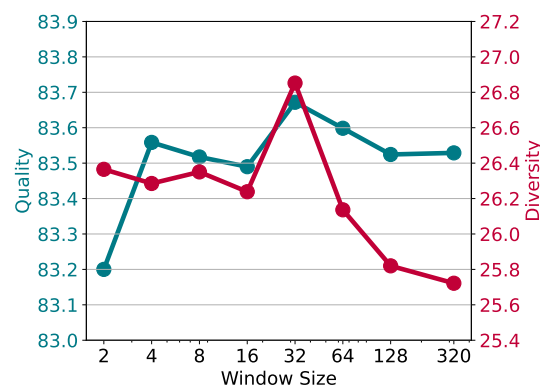


Figure 1: Quality and diversity of generated text with varying attention sparsity on ROCStories. The quality and diversity (the higher the better) are measured by BERTScore (Zhang et al., 2020) and dist (Li et al., 2016a), respectively. Sparsity is controlled by the attention window size. A smaller window leads to higher sparsity.

tend to remember and generate *dull* and *generic* (usually high-frequency) contents (Holtzman et al., 2020), even with highly-diverse ground-truth as training targets and distinct prompts as testing inputs. Such a conundrum, known as one kind of *degeneration problem* (Holtzman et al., 2018) in NLG, could cause a poor user experience, especially for genres with high requirements for diversity and novelty, like story and headline generation. Several research lines make continued endeavors for this challenge, for instance, adopting stochastic decoding algorithms (Kool et al., 2019; Holtzman et al., 2020) or replacing the Maximum Likelihood Estimation (MLE) loss with more sophisticated optimization objectives (Welleck et al., 2020; Xu et al., 2022a) to alleviate degeneration. However, these methods are just stopgaps and do not address this issue in depth. The stochastic decoding methods permit the model to choose low-confidence candidates, therefore bringing the contradiction between the quality and the diversity of the generated text. Meanwhile, the methods in the second

*Corresponding author. Email: sms@tsinghua.edu.cn

research line just focus on solving *intra-instance repetition* instead of *inter-instance diversity*. Thus, to intrinsically improve the diversity in text generation without hurting the quality, the underlying mechanism of this problem must be figured out.

What leads to poor generation diversity and how could we accordingly tackle it? In this work, we dig into these questions and find another possible cause of such a conundrum in Transformer-based models, namely *Attention Concentration*. Initially inspired by related literary research which demonstrated that the concentration degree of attention on linguistic sources could impact creative writing (Rosenblatt, 1988), we conducted a preliminary study on the correlation between attention concentration (or sparsity) and generation diversity. We trained a GPT-2 base model with local attention window as in Child et al. (2019) on the ROCStories dataset (Mostafazadeh et al., 2016). As shown in Fig. 1, obviously, adequate sparsity improves both quality and diversity, while excessive sparsity harms quality. Such improvement of diversity is non-trivial as demonstrated in Sec. 4.

To better understand such results, we resort to the representation degeneration problem (Gao et al., 2019) and analyze the interaction between the hidden states and word embeddings in Transformer. We discovered that the attentive mixture could impel hidden states to approach the embeddings of high-frequency words, especially when the attention is highly distributed, which encourages irrationally higher probabilities of the generic high-frequency words, harming diversity and novelty. We termed this phenomenon as the *Trap of Mediocrity*.

To escape this trap enforced by evenly dispersed attention, we propose a simple yet effective regularization method, CARE¹, to concentrate and sharpen attention. CARE reduces the Rényi entropy of attention distributions during the training process to enhance sparsity, which is lightweight, transparent to the model architectures, and could be easily implemented within 20 lines of code. We mathematically demonstrate CARE as learning a Bayesian approximation of the posterior attention given the ground truth, providing a theoretical validation of its superiority. Besides, We also equip CARE with learnable attention dropout to further boost performance. In this way, CARE could learn more concentrated attention with minimal computa-

tion cost to break the Trap of Mediocrity and hence improve the diversity and novelty of generated text.

In summary, our contributions are as follows:

- We are the first work to find the correlation between attention concentration and generation diversity, *i.e.*, *Trap of Mediocrity*, and provide an in-depth analysis of its underlying cause.
- We propose a lightweight method, CARE, to escape from such a trap, which is transparent to model structures and easy to implement.
- We theoretically demonstrate that CARE can concentrate attention and learn a Bayesian approximation of the posterior attention.
- Experiments on three conditional and one unconditional generation tasks show that our model could significantly improve generation diversity and novelty against several strong baselines while keeping comparable quality.

2 The Trap of Mediocrity

As introduced in Sec. 1, we find one underlying cause of poor generation diversity lies in the improperly attentive mixture of the hidden states in Transformer, *i.e.*, *Trap of Mediocrity*. To reveal the mechanism of this trap, we performed in-depth analyses both empirically and theoretically. In detail, we progressed our argumentation by proving the following three **Research Claims (RC)** step by step:

RC1: Formation of the trap. The output-layer embeddings of high-frequency tokens are more likely to cluster in a uniformly positive direction of corresponding hidden states in the Transformer.

RC2: Falling into the trap The attentive mixture process of contextual hidden states forces the representations of all tokens to approach that direction, resulting in consistently higher generation probabilities of frequent and mediocre tokens.

RC3: Dispersion dominates the mixture The extremely dispersed attention distribution (low sparsity) loses the concentration on informative context and dominates the improper mixture both in the formation and falling process.

To support each of these claims, we reuse the GPT-2 in Sec. 1 and refer to the hidden state at the top Transformer layer (before the output softmax) as *hidden state* for brevity in this section².

¹Concentrating Attention by Rényi Entropy

²We also similarly analyzed unconditional NLG on the

2.1 RC1: Formation of the Trap

We decompose RC1 into three sub-claims and successively verify each of them.

Sub-claim 1: The output-layer embeddings of frequent tokens are more likely to approach corresponding hidden states during training. Define the output embedding vectors as $\{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_V\}$ with vocabulary size V and $\mathbf{w}_k \in \mathbb{R}^{1 \times d}$ corresponding to the k -th token x_k where d is hidden size, c as the input condition (prompt). Following (Gao et al., 2019), we consider the interaction of each word embedding and hidden state. During typical training, we aim to maximize the likelihood of the target token x_i at the t -th time step:

$$P(y_t = x_i | c, y_{<t}) = \frac{\exp(\mathbf{w}_i \mathbf{h}_t^T)}{\exp(\mathbf{w}_i \mathbf{h}_t^T) + C} = 1 - \frac{C}{\exp(\mathbf{w}_i \mathbf{h}_t^T) + C}, \quad (1)$$

where \mathbf{h}_t is the hidden state at t -th time step, $y_{<t}$ means the preceding $t-1$ tokens, and $C = \sum_{j, j \neq i} \exp(\mathbf{w}_j \mathbf{h}_t^T)$. By ignoring the terms irrelevant to \mathbf{w}_i , i.e., C , we can attribute the loss mainly to $\mathbf{w}_i \mathbf{h}_t^T$ and further factorize it as:

$$\mathbf{w}_i \mathbf{h}_t^T = \|\mathbf{w}_i\| \|\mathbf{h}_t\| \cos \theta. \quad (2)$$

From Eq. (2), we could derive that maximizing Eq. (1) actually minimizes the angle θ between \mathbf{w}_i and \mathbf{h}_t . Since $\|\mathbf{h}_t\|$ is bounded by the layer normalization (Xiong et al., 2020), smaller $\|\mathbf{w}_i\|$ would require smaller θ to reach the same low loss. Based on this conclusion, we further observed that embeddings of high-frequency tokens tend to possess smaller norms, as shown in Fig. 2. As a result, we successfully verified sub-claim 1.

Sub-claim 2: The embeddings of high-frequency tokens tend to cluster. We investigated the aggregation degree of token embedding, as shown in Fig. 3. It is obvious that rare tokens are clustered, as reported in (Yu et al., 2022a). Furthermore, we can find that the highly frequent (top 10%) tokens are also closer to each other than the mid-frequency ones, which is also observed in BERT (Li et al., 2020a), empirically supporting sub-claim 2.

Sub-claim 3: There is a uniformly positive direction of the hidden states corresponding to the frequent tokens, i.e., $\mathbf{w}_i \mathbf{h}_t^T > 0$. By slightly modifying the conclusion in (Gao et al., 2019), we could

Yelp dataset (Zhang et al., 2015) and got the same conclusion. We leave the details in Appendix B due to the space limit.

demonstrate that there exists such a direction. See Appendix C for detailed derivation.

By combining these sub-claims, we could infer that embeddings of high-frequency tend to cluster and approach a uniformly positive direction of the hidden states during training, manifesting RC1.

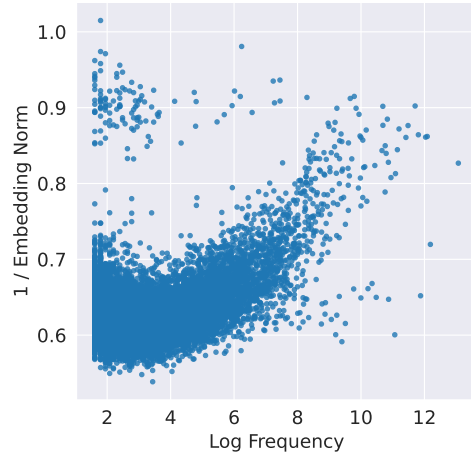


Figure 2: The correlation between reciprocal embedding norm and logarithmic word frequency. The salient positive correlation ($R^2 = 0.835$) indicates that high-frequency words tend to have smaller embedding norms.

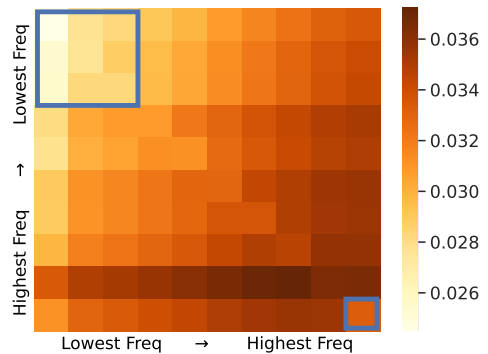


Figure 3: The distance of embeddings in/between different frequency intervals. We divided token embeddings uniformly into 10 intervals by frequency and calculated the average l_2 distance between each embedding and the top-50 nearest ones in each interval.

2.2 RC2: Falling into the Trap

We have demonstrated how the trap is formed, which underlines the mediocre generated text. In this subsection, we further explore how the model falls into such a trap during the generation process.

In the Transformer self-attention, the hidden

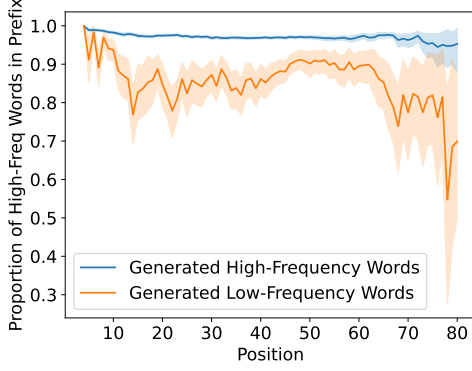


Figure 4: The proportion of high-frequency words in the current sub-sequence when a high/low-frequency token is generated. The gap between the two lines indicates that low-frequency words tend to be generated from the context containing less high-frequency words. The sequences were generated by beam search.

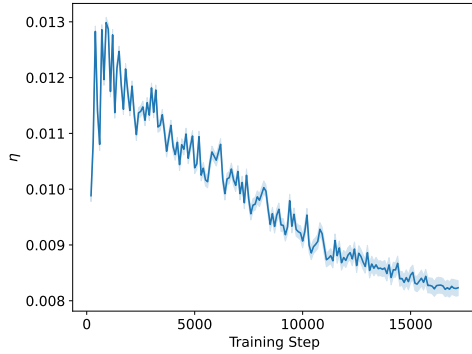


Figure 5: The value of η in Sub-claim 4 during training.

states H^l in l -th layer are calculated as:

$$\mathbf{h}_t^l = \text{softmax} \left(\frac{\mathbf{h}_t^{l-1} \mathbf{W}_q^l (\mathbf{H}^{l-1} \mathbf{W}_k^l)^T}{\sqrt{d}} \right) \mathbf{H}^{l-1} \mathbf{W}_v^l, \quad (3)$$

where $\mathbf{H}^{l-1} = [\mathbf{h}_1^{l-1}; \dots; \mathbf{h}_T^{l-1}]$ are hidden states of the previous layer and $\mathbf{W}_q^l, \mathbf{W}_k^l, \mathbf{W}_v^l$ are query, key and value projection matrix, respectively.

From Eq. (3), we can clearly see that each \mathbf{h}_t would mix multiple previous hidden states by the attention score. We name such mixture as the *attentive mixture*. Though in RC1 we only focus on hidden states corresponding to frequent target tokens, these high-frequency tokens are more likely to appear in the context of previous tokens, which could push the mixed \mathbf{h}_t towards the uniformly positive direction and hence draw it closer to the cluster of high-frequency embeddings. As a result, the model falls into *the Trap of Medicority*.

Moreover, as the sequence becomes longer, the

number of high-frequency tokens also increases, which further aggravates the tendency of frequent tokens and keeps trapping the model. Fig. 4 manifests such deterioration. We can observe that with a longer sentence, GPT-2 requires increasing more infrequent tokens in context to motivate a low-frequency token. Inversely, the context with a large proportion of high-frequency tokens would always encourage frequent tokens at the next time step. Besides, we also calculated the pairwise cosine similarity between embeddings and the hidden states of all sentences in the testing set and then found that compared to the low-frequency embeddings, the high-frequency ones are generally nearer to all hidden states by a stable and large margin.

By these observations, we can empirically conclude that the attentive mixture process of hidden states forces a closer distance between each hidden state and frequent token embeddings, and hence causes higher generation probabilities (see Eq. (1)) of high-frequency words, validating RC2.

2.3 RC3: Dispersion dominates the mixture

Subsections 2.1 & 2.2 describe the complete process of how the trap of mediocrity forms and how the model falls into it. This process is dominated by the attentive mixture and would further deteriorate with a more scattered attention distribution.

To theoretically verify this claim, following (Yu et al., 2022a), we investigate the negative log-likelihood loss of a target sequence with length T :

$$\mathcal{L}_{NLL} = -\frac{1}{T} \sum_{t=1}^T \log \frac{\exp(\mathbf{w}_{x_t} \mathbf{h}_t^T)}{\sum_j^V \exp(\mathbf{w}_j \mathbf{h}_t^T)}, \quad (4)$$

where x_t is the ground-truth token at the t -th step.

We rewrite the attention score used in Eq. (3) as:

$$\mathbf{a}_t = \text{softmax} \left(\frac{\mathbf{h}_t \mathbf{W}_q^l (\mathbf{H} \mathbf{W}_k^l)^T}{\sqrt{d}} \right), \quad (5)$$

where $\mathbf{a}_t = [a_{t,1}, \dots, a_{t,t}]$ represent the attention weights towards the preceding tokens.

For simplification, we ignore the nonlinearity in attention layers and let $\mathbf{h}_t^l = \sum_j a_{t,j} \hat{\mathbf{h}}_j$, where $\hat{\mathbf{h}}_j = \mathbf{h}_j^{l-1} \mathbf{W}_v^l$, and give the following conclusion:

Sub-claim 4: With the extremely dispersed attention distribution which reaches the highest entropy, i.e., for each time step t , $a_{t,j} = \frac{1}{t}, j \in \{1, \dots, t\}$, the optimization direction of the embeddings would approach the uniformly positive direction of all hidden states if $\eta = \frac{1}{t} \sum_{j=1}^t \frac{p_j}{j} > 0$ for each t , where

p_j is the generation probability of the ground-truth token x_i of the t -th step, $p_t = p(y_t = x_i | c, y_{<t})$.

A theoretical demonstration of this sub-claim could be referred to Appendix C.

Sub-claim 4 indicates that when $\eta > 0$, the improperly scattered attention would push embeddings towards the uniformly positive direction where high-frequency embeddings cluster, rather than more diverse directions, which is fundamental in forming the trap as described in Sec. 2.2. We further examine η during the training process and plot the result in Fig. 5. Generally, η is positive during the whole training process, proving RC3.

By combining RCs 1~3, we successfully demonstrate that the improperly scattered attention is the underlying cause of poor diversity. A proper attention score should concentrate on some specific parts of the context, while in practice, the model would lose concentration and fall into the *Trap of Medicority*, which implies we could break the trap and improve diversity by concentrating attention.

3 Methodology

3.1 CARE

According to our analyses in Sec. 2, the improperly dispersed (low-sparsity) attention is one fundamental cause of the *Trap of Medicority*. Therefore, we propose a novel regularization method, CARE, to concentrate attention and thus promote diversity.

In detail, we add a regularization term to the original Maximum Likelihood Estimate (MLE) loss:

$$\mathcal{L}_R = \sum_{t=1}^T \frac{\alpha(t+1)}{t(\alpha-1)} \|\hat{\mathbf{a}}_t\|_1, \quad (6)$$

$$\mathcal{L}_{\text{CARE}} = \mathcal{L}_{\text{MLE}} + \gamma \mathcal{L}_R, \quad (7)$$

where $\hat{\mathbf{a}}_t$ denotes the unnormalized attention logits before softmax and attention dropout, $\|\cdot\|_1$ is the l_1 norm, T is the sequence length, $\alpha > 1$ is a hyperparameter to control the sharpness of the attention distribution, and γ could be utilized to balance the MLE term and the regularization term.

To theoretically demonstrate our method’s effectiveness in enhancing attention sparsity, we present the following conclusion:

Theorem 1 *Minimizing \mathcal{L}_R in Eq. (6) is equal to minimizing an upper bounds of the Rényi entropy*³

³https://en.wikipedia.org/wiki/R%C3%A9nyi_entropy

of the attention distribution $H_\alpha(\mathbf{a}_t)$, where H_α denotes the Rényi Entropy and α is the sharpness controller in Eq. (6).

Theorem 1 reveals that by minimizing \mathcal{L}_R , we could reduce the entropy of the attention distribution and hence make attention more concentrated, which helps evade the trap and boost diversity. Please refer to Appendix C for detailed proof.

Furthermore, \mathcal{L}_R is self-adaptive benefiting from the derived weight $\beta = \frac{\alpha(t+1)}{t(\alpha-1)}$. As shown in Fig. 4, longer generated sequences (larger t) exacerbate the trap, but larger t also leads to a smaller β , which forces the model to reduce $\|\hat{\mathbf{a}}_t\|_1$ more to reach the same low loss and vice versa. This adaptability could mitigate the deterioration caused by length.

Besides, we could also consider the benefits of CARE from the perspective of *Bayesian Inference*. Following (Gal and Ghahramani, 2016), we give another conclusion to validate such an advantage:

Theorem 2 *When combined with the attention dropout, optimizing $\mathcal{L}_{\text{CARE}}$ can be regarded as minimizing $KL[q_\theta(\tilde{\mathbf{a}}_t) || p(\tilde{\mathbf{a}}_t | c, x_t, y_t)]$, that is, learning a Bayesian approximation $q_\theta(\tilde{\mathbf{a}}_t)$ of the true posterior attention $p(\tilde{\mathbf{a}}_t | c, x_t, y_t)$ at each time step, where x_t, y_t denote the input and target tokens at the t -th step, respectively, and $\tilde{\mathbf{a}}_t$ denotes the unnormalized attention logits before softmax and after attention dropout,*

Proof. See Appendix C.

Theorem 2 demonstrates another strong endorsement of our model: CARE could act as a posterior estimator of the attention distribution, which helps improve the accuracy of attention and hence maintain comparable generation quality. In comparison, straightforward sparse attention methods without such theoretical guarantees can improve diversity but also hurt quality. See Table 1 for detailed results.

In addition, to conform to this theorem, we need to apply attention dropout to the unnormalized attention distribution, achieved by by sampling the dropping mask from a Bernoulli distribution and accordingly adding a large negative constant to the unnormalized attention logits. We also compare this setting with the the original one in Transformer in Appendix D and manifested the superiority of ours.

Noth that CARE is transparent to different model architectures since it involves only the output attention weights in Transformer, and hence can be

easily implemented within 20 lines of codes. We provide the implementation details in Appendix E.

3.2 CARE-A

Since Theorem 2 indicates that the attention dropout is crucial for our method, we could further enhance the capacity of CARE by incorporating the more flexible Concrete Dropout (Gal et al., 2017). The original dropout can be considered as sampling attention masks z from a Bernoulli distribution $z \sim B(1, p)$ which could influence attention concentration through a specified dropout ratio p . In contrast, Concrete Dropout relaxes z as $\tilde{z} = \sigma((\log p - \log(1-p) + \log u - \log(1-u))/\tau)$, where p is learnable, τ is the temperature, σ means sigmoid, and $u \sim \text{Uniform}(0, 1)$. As the optimization requires an extra entropy term of p as described in (Gal et al., 2017), we rewrite $\mathcal{L}_{\text{CARE}}$ as:

$$\mathcal{L}_{\text{CARE-A}} = \mathcal{L}_{\text{MLE}} + \gamma \mathcal{L}_R + \delta \sum_{i=1}^L H(p_i), \quad (8)$$

where L is the number of Transformer layers.

In this way, we learn layer-wise and adaptive dropout ratios p , benefiting a more flexible control of attention concentration. Therefore, we name this variant of our model as CARE-A.

4 Experiment

4.1 Datasets

We conducted comprehensive experiments on three conditional generation tasks: story generation on ROCStories (Mostafazadeh et al., 2016), headline generation on MIND (Wu et al., 2020) and PENS (Ao et al., 2021), and paraphrase generation on ParaSCI (Dong et al., 2021a). Specifically, we use the ArXiv set in ParaSCI which contains more instances. MIND and PENS are merged and re-split for more training data. For unconditional generation, we utilized the popular Yelp dataset (Zhang et al., 2015). We truncated the text in these datasets to a max length of 1024 tokens. Detailed data statistics are listed in Appendix A.

4.2 Implementation Details

We use GPT-2 base (Radford et al., 2019) as the backbone for ROCStories, ParaSCI, and Yelp, and a little smaller model for headline generation due to the limit of our computational power. We train our CARE, all baseline models, and the BPE tokenizers from scratch on each dataset. We use

beam search for conditional generation datasets to avoid the diversity from sampling randomness, and sampling decoding for unconditional generation since beam search may bring too much duplicated content on Yelp. All models are implemented by HuggingFace library⁴ and are trained from scratch. See Appendix A for more setting details.

4.3 Baselines

Besides the original (a) **GPT-2** (Radford et al., 2019), since our main claim is the correlation between attention concentration and generation diversity, we also compared several strong sparse attention based methods. (b) **Pattern**: the fixed attention pattern with local and sliding windows (Child et al., 2019). (c) **Entmax** (Correia et al., 2019): a learned adaptive sparse attention method. (d) **l_0 -Drop** (Zhang et al., 2021): this method uses attention masks sampled from the hard concrete distribution to achieve sparsity and control sparsity degree by the l_0 norm loss on the mask. (e) **\mathcal{L}_A -Tuning** (Manakul and Gales, 2021): a method which directly adds an entropy term of attention distributions to the training loss. All models share the same configuration for fair comparisons. More details of baseline models are provided in Appendix A.

4.4 Metrics

For the conditional generation tasks, we evaluate the **quality** of the generated text by BLEU (Papineni et al., 2002), ROUGE (Lin and Hovy, 2002), BertScore (Zhang et al., 2020), and CND (Li et al., 2020b). We report BLEU-2,4 and ROUGE-2,3,L,W, which are most commonly used in the literature. For **diversity**, we assess Dist (Li et al., 2016a), JS (Wang and Wan, 2018), and Self-BLEU (Zhu et al., 2018). We report the geometric mean of 1-gram to 4-gram for CND, Dist, JS, and Self-BLEU. For unconditional generation, we report BLEU, MAUVE (Pillutla et al., 2021), and CND for quality measurement. MAUVE is multiplied by 100 for convenience. For diversity, we take the same three metrics in the conditional generation.

4.5 Results

4.5.1 Conditional Generation

Table 1 presents the evaluation results on ROCStories and ParaSCI. We leave those of headline

⁴<https://huggingface.co/>

Model	Quality								Diversity		
	R-2 \uparrow	R-3 \uparrow	R-L \uparrow	R-W \uparrow	B-2 \uparrow	B-4 \uparrow	BS \uparrow	CND \downarrow	Dist \uparrow	JS \downarrow	SB \downarrow
Dataset: ParaSCI											
GPT-2	41.48	32.65	54.97	35.28	41.64	<u>26.58</u>	90.81	1.688	60.34	0.0726	17.52
Pattern	41.46	32.58	55.06	35.34	41.31	26.43	<u>90.89</u>	1.652	60.65	0.0683	16.86
Entmax	38.44	29.78	52.28	33.50	38.06	23.60	90.35	1.769	59.30	0.0707	18.10
l_0 -Drop	37.26	28.64	50.99	32.53	37.96	23.47	90.05	1.700	60.94	<u>0.0591</u>	<u>16.18</u>
\mathcal{L}_A -Tuning	40.71	32.05	54.20	34.89	39.95	25.48	90.70	<u>1.632</u>	60.67	0.0622	16.28
CARE	42.49	33.74	<u>55.69</u>	35.94	41.59	26.75	90.80	1.631	61.04	0.0566	16.05
CARE-A	<u>41.95</u>	<u>33.01</u>	55.79	<u>35.76</u>	<u>41.63</u>	26.56	91.04	1.637	60.60	0.0596	17.04
Dataset: ROCStories											
GPT-2	6.700	1.321	23.17	<u>11.19</u>	5.817	0.266	83.53	8.856	25.72	0.6625	62.36
Pattern	6.690	1.301	23.33	11.09	6.207	0.273	<u>83.64</u>	9.006	27.20	0.6280	61.16
Entmax	6.011	1.011	22.42	10.80	5.374	0.150	82.91	9.521	25.03	0.7317	62.57
l_0 -Drop	5.850	1.065	22.11	10.77	4.896	0.177	83.17	8.489	24.82	0.6207	62.72
\mathcal{L}_A -Tuning	6.367	1.214	22.94	10.96	5.809	0.226	83.37	8.904	<u>27.24</u>	<u>0.5905</u>	<u>61.06</u>
CARE	<u>6.905</u>	1.399	<u>23.54</u>	11.08	6.829	0.326	83.74	9.735	27.77	0.6328	60.41
CARE-A	7.002	<u>1.332</u>	23.80	11.32	<u>6.489</u>	<u>0.274</u>	83.63	<u>8.742</u>	25.89	0.5708	61.51

Table 1: Evaluation results for conditional generation. The best results are **bolded** and the second best ones are underlined. B, R, SB and BS represent BLEU, ROUGE, Self-BLEU and BertScore, respectively.

generation in Appendix D due to space limitations. We can see that compared to the naive GPT-2, most models with sparse attention improve generation diversity to some extent, except for Entmax which is relatively unstable due to its iterative estimation of the Tsallis entropy. Such results could well support our motivation and claim that the scattered attention causes the *Trap of Mediocrity* and more concentrated attention helps evade the trap.

However, these sparse attention baseline methods significantly decrease the quality of generated text when optimizing diversity, while CARE(-A) attained superior performance in novelty and diversity and meanwhile even boosted the quality of the generated text, demonstrating its effectiveness from Theorem 2. In addition, Pattern also increased the diversity without hurting quality, but is still inferior to our method under most metrics.

4.5.2 Unconditional Generation

Model	Quality				Diversity		
	B-2 \uparrow	B-4 \uparrow	CND \downarrow	MV \uparrow	Dist \uparrow	JS \downarrow	SB \downarrow
GPT-2	76.68	36.46	<u>0.886</u>	<u>0.931</u>	27.62	0.187	47.76
Pattern	76.96	<u>38.09</u>	0.924	0.859	27.34	0.203	49.31
Entmax	77.03	32.18	0.951	0.920	23.99	0.222	49.18
l_0 -Drop	76.71	34.92	0.866	0.919	26.33	0.194	47.38
\mathcal{L}_A -Tuning	<u>77.35</u>	38.66	0.910	0.835	26.26	0.211	50.82
CARE	73.84	33.61	1.033	0.863	30.84	<u>0.178</u>	43.60
CARE-A	77.53	35.69	0.933	0.974	<u>28.44</u>	0.169	<u>44.65</u>

Table 2: Evaluation results for unconditional generation on the Yelp dataset. MV stands for MAUVE.

The results of unconditional generation are pre-

sented in Table 2. In this task, most sparse baselines perform no better than GPT-2 on diversity. This phenomenon could be attributed to the sufficient diversity brought by the sampling decoding. However, our method can still generate more diverse content compared to GPT-2, manifesting its extreme power in evading mediocrity and enhancing novelty. CARE-A even achieves comparable quality while significantly improving diversity.

4.6 Human Evaluation

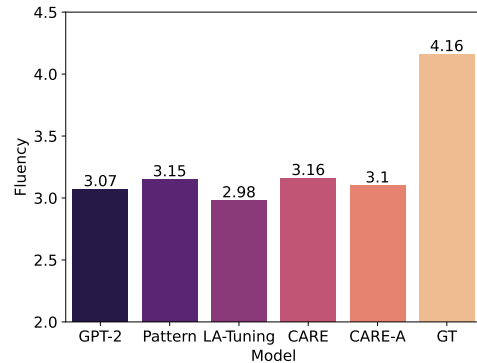


Figure 6: Human evaluation on the fluency of the generated text. GT stands for human-written ground truth. The p-value < 0.001 , and the Kappa score is 0.67, indicating an acceptable inter-annotator agreement.

We also conduct a human evaluation of the fluency of generated texts on ROCStories. See Appendix A.5 for the concrete evaluation protocol. As shown in Table 6, the fluency of the text gener-

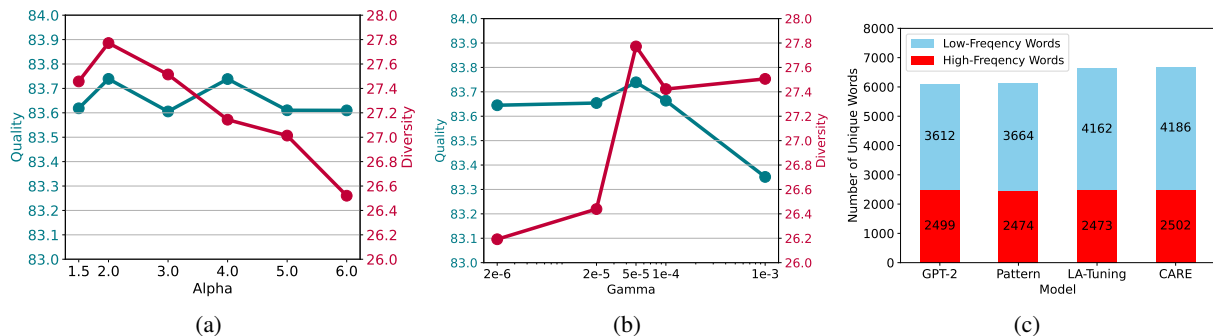


Figure 7: (a, b) The performance when varying different α and γ . The quality of the generated text is measured by BERTScore and diversity by Dist. (c) The unique high-frequency words in the generated text of each model.

ated by CARE is comparable with GPT-2 and other baseline methods. Combined with the automatic evaluation, we could conclude that our model obtains better diversity without losing any fluency.

4.7 Ablation Study

To further investigate the effect of \mathcal{L}_R , we vary different values of α and γ on ROCStories and present the results in Fig. 7. Fig. 7a shows that our model benefits from the adequate sharpness (larger α) by the entropy regularization, while excessively sharp or blunt attention hurts both quality and diversity. Besides, the weight of \mathcal{L}_R also needs to be carefully tuned since an overly small weight cannot exert the effect on enhancing sparsity. In contrast, an unduly large weight (too large γ) would hinder the optimization of MLE and result in a decline in generation quality, as shown in Fig. 7b.

4.8 Novelty Analysis

To verify whether our method achieves better novelty by inclining to the low-frequency words, we separately count the unique high/low-frequency words generated by CARE in ROCStories following (Yu et al., 2022b). Please note the subtle difference between diversity and novelty. Only sufficient novel (distinct) words can support satisfactorily good inter-instance diversity. As presented in Fig. 7c, the result demonstrates that our model generates more unique low-frequency words than GPT-2 and Pattern by a large margin and \mathcal{L}_A Tuning by a small margin, which indicates that the diversity and novelty brought by sparse attention could be attributed to the incline of low-frequency words. This finding again confirms that the diversity improvement comes from evading the *Trap of Mediocrity*, while our method is the most successful Escapist.

4.9 Case Study

Prompt: i went to the park to play volleyball .
GPT-2: i met up with a few people. we played a few games of pool volleyball. it was a lot of fun. i had a great time playing volleyball.
Pattern: it was my turn to serve the ball. i got the ball. i threw the ball to the receiver. the ball was intercepted.
CARE: i met up with a lot of my friends there and we split into two teams. each team had nine players on it. [FEMALE] team scored first with a home run. [FEMALE] team won the game in convincing fashion.

Figure 8: A generated sample in ROC. Generic / repetitive continuations are marked in different colors, while those in pink / red are relevant/novel expressions.

Fig. 8 shows a typical sample generated by CARE and the baselines. GPT-2 generates relatively generic continuations with mediocre descriptions like *it was a lot of fun* and simply copies the word *volleyball* several times, while models with sparse attention could generate novel and rich content precise to the given topic. CARE could further generate more novel and topic-related phrases like *home run* and *convincing fashion*.

5 Related Work

Enhancing Diversity in NLG Previous work focusing on diversity in text generation mainly falls into two lines. The first line eschews this issue by incorporating randomization and diversification into decoding algorithms, e.g., top-k (Fan et al., 2018) and top-p (Holtzman et al., 2020) sampling, diversified beam search (Li et al., 2016b; Vijayakumar et al., 2018), and stochastic beam search (Kool et al., 2019). The other line aims to substitute or supplement the MLE loss with novel objectives, which involves Reinforcement Learning (Shi et al., 2018), VAE (Nie et al., 2019; Hu et al., 2022), ad-

versarial training (Diao et al., 2021; Zhou et al., 2021), unlikelihood training (Welleck et al., 2020), energy-based models (Deng et al., 2020), and additional penalty terms (Su et al., 2022; Xu et al., 2022b). There are also works that modify attention to handle this problem (Sun et al., 2020; Dong et al., 2021b). However, they only consider on-the-fly modification instead of analyzing the comprehensive relationship between attention and NLG diversity.

Sparse Attention There have been continuous efforts to reduce the computation cost of self-attention in Transformer. The majority make the dense attention matrix sparser to improve efficiency, including fixed patterns (Child et al., 2019; Beltagy et al., 2020; Ainslie et al., 2020; Zaheer et al., 2020; Sukhbaatar et al., 2019; Huang et al., 2021), cluster-based methods (Wang et al., 2020; Roy et al., 2021), learnable patterns (Correia et al., 2019; Kitaev et al., 2020; Tay et al., 2020a), and regularization-based methods (Zhang et al., 2021; Shi et al., 2021). More details can be referred to (Tay et al., 2020b). However, all these methods focus on saving computational budget and improving the performance of language modeling. We are the first to study sparse attention from the in-depth view of improving generation diversity.

6 Conclusion

In this paper, we analyze the underlying cause of generic generated texts and propose a theory called the *Trap of Mediocrity* to attribute this problem to the attention dispersion in Transformer. To evade such a trap, we present CARE to concentrate attention, which is easy to implement and transparent to model structures. Experiments on a variety of NLG tasks demonstrate CARE’s superiority in improving diversity while maintaining competitive quality. The code of our method will be released in <https://github.com/peterliwenhao/CARE>.

Acknowledgement

We sincerely thank all the anonymous reviewers for their valuable comments. This work is supported by the National Key R&D Program of China (No. 2020AAA0106502) and Institute Guo Qiang at Tsinghua University.

Limitation

There are still several limitations to our methods. Firstly, due to the introduction of new hyperparameters, our method may need extra hyperparameter tuning when adapted to the specific task. Secondly, the validity of our approach to finetuning models pretrained with full attention has not been tested, which needs further experiments. Since we are the first to propose the *Trap of Mediocrity* to our best knowledge, our exploration of this problem is inevitably very preliminary or even defective. We will keep refining our methods and expect more future work to further explore this problem.

References

- Joshua Ainslie, Santiago Ontañón, Chris Alberti, Václav Cvicek, Zachary Fisher, Philip Pham, Anirudh Ravula, Sumit Sanghai, Qifan Wang, and Li Yang. 2020. **ETC: encoding long and structured inputs in transformers**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 268–284. Association for Computational Linguistics.
- Xiang Ao, Xiting Wang, Ling Luo, Ying Qiao, Qing He, and Xing Xie. 2021. **PENS: A dataset and generic framework for personalized news headline generation**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 82–92. Association for Computational Linguistics.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. **Longformer: The long-document transformer**. *CoRR*, abs/2004.05150.
- Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. **Generating long sequences with sparse transformers**. *arXiv preprint arXiv:1904.10509*.
- Gonçalo M. Correia, Vlad Niculae, and André F. T. Martins. 2019. **Adaptively sparse transformers**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2174–2184, Hong Kong, China. Association for Computational Linguistics.
- Yuntian Deng, Anton Bakhtin, Myle Ott, Arthur Szlam, and Marc’Aurelio Ranzato. 2020. **Residual energy-based models for text generation**. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

- Shizhe Diao, Xinwei Shen, Kashun Shum, Yan Song, and Tong Zhang. 2021. [TILGAN: transformer-based implicit latent GAN for diverse and coherent text generation](#). In *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, volume ACL/IJCNLP 2021 of *Findings of ACL*, pages 4844–4858. Association for Computational Linguistics.
- Qingxiu Dong, Xiaojun Wan, and Yue Cao. 2021a. [Parasci: A large scientific paraphrase dataset for longer paraphrase generation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, pages 424–434. Association for Computational Linguistics.
- Yue Dong, Chandra Bhagavatula, Ximing Lu, Jena D. Hwang, Antoine Bosselut, Jackie Chi Kit Cheung, and Yejin Choi. 2021b. [On-the-fly attention modulation for neural generation](#). In *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, volume ACL/IJCNLP 2021 of *Findings of ACL*, pages 1261–1274. Association for Computational Linguistics.
- Angela Fan, Mike Lewis, and Yann N. Dauphin. 2018. [Hierarchical neural story generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 889–898. Association for Computational Linguistics.
- Yarin Gal and Zoubin Ghahramani. 2016. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR.
- Yarin Gal, Jiri Hron, and Alex Kendall. 2017. [Concrete dropout](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 3581–3590.
- Jun Gao, Di He, Xu Tan, Tao Qin, Liwei Wang, and Tie-Yan Liu. 2019. [Representation degeneration problem in training natural language generation models](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Jian Guan, Fei Huang, Zhihao Zhao, Xiaoyan Zhu, and Minlie Huang. 2020. [A Knowledge-Enhanced Pre-training Model for Commonsense Story Generation](#). *Transactions of the Association for Computational Linguistics*, 8:93–108.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text degeneration](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Ari Holtzman, Jan Buys, Maxwell Forbes, Antoine Bosselut, David Golub, and Yejin Choi. 2018. [Learning to write with cooperative discriminators](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 1638–1649. Association for Computational Linguistics.
- Jinyi Hu, Xiaoyuan Yi, Wenhao Li, Maosong Sun, and Xing Xie. 2022. [Fuse it more deeply! a variational transformer with layer-wise latent variable inference for text generation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 697–716, Seattle, United States. Association for Computational Linguistics.
- Luyang Huang, Shuyang Cao, Nikolaus Nova Parulian, Heng Ji, and Lu Wang. 2021. [Efficient attentions for long document summarization](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 1419–1436. Association for Computational Linguistics.
- Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. 2020. [Reformer: The efficient transformer](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Wouter Kool, Herke van Hoof, and Max Welling. 2019. [Stochastic beams and where to find them: The gumbel-top-k trick for sampling sequences without replacement](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 3499–3508. PMLR.
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario Sasko, Guntan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue, Théo Matussière, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François Lagunas, Alexander M. Rush, and Thomas Wolf. 2021. [Datasets: A community library for natural language processing](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, EMNLP 2021, Online and Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 175–184. Association for Computational Linguistics.
- Bohan Li, Hao Zhou, Junxian He, Mingxuan Wang, Yiming Yang, and Lei Li. 2020a. [On the sentence embeddings from pre-trained language models](#). In *Proceedings of the 2020 Conference on Empirical*

- Methods in Natural Language Processing (EMNLP)*, pages 9119–9130, Online. Association for Computational Linguistics.
- Jianing Li, Yanyan Lan, Jiafeng Guo, and Xueqi Cheng. 2020b. On the relation between quality-diversity evaluation and distribution-fitting goal in text generation. In *International Conference on Machine Learning*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016a. [A diversity-promoting objective function for neural conversation models](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2016b. [A simple, fast diverse decoding algorithm for neural generation](#). *CoRR*, abs/1611.08562.
- Chin-Yew Lin and Eduard Hovy. 2002. [Manual and automatic evaluation of summaries](#). In *Proceedings of the ACL-02 Workshop on Automatic Summarization*, pages 45–51, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Zhe Lin, Yitao Cai, and Xiaojun Wan. 2021. [Towards document-level paraphrase generation with sentence rewriting and reordering](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021*, pages 1033–1044. Association for Computational Linguistics.
- Potsawee Manakul and Mark J. F. Gales. 2021. [Sparsity and sentence structure in encoder-decoder attention of summarization systems](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 9359–9368. Association for Computational Linguistics.
- Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. 2016. [A corpus and cloze evaluation for deeper understanding of commonsense stories](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 839–849, San Diego, California. Association for Computational Linguistics.
- Weili Nie, Nina Narodytska, and Ankit Patel. 2019. [Relgan: Relational generative adversarial networks for text generation](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Krishna Pillutla, Swabha Swayamdipta, Rowan Zellers, John Thickstun, Sean Welleck, Yejin Choi, and Zaid Harchaoui. 2021. Mauve: Measuring the gap between neural text and human text using divergence frontiers. In *Neural Information Processing Systems*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Louise Michelle Rosenblatt. 1988. *Writing and reading: The transactional theory*. 416. University of Illinois at Urbana-Champaign.
- Aurko Roy, Mohammad Saffar, Ashish Vaswani, and David Grangier. 2021. [Efficient content-based sparse attention with routing transformers](#). *Trans. Assoc. Comput. Linguistics*, 9:53–68.
- Han Shi, Jiahui Gao, Xiaozhe Ren, Hang Xu, Xiaodan Liang, Zhenguo Li, and James Tin-Yau Kwok. 2021. [Sparsebert: Rethinking the importance analysis in self-attention](#). In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 9547–9557. PMLR.
- Zhan Shi, Xinchu Chen, Xipeng Qiu, and Xuanjing Huang. 2018. [Toward diverse text generation with inverse reinforcement learning](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 4361–4367. ijcai.org.
- Yixuan Su, Tian Lan, Yan Wang, Dani Yogatama, Lingpeng Kong, and Nigel Collier. 2022. [A contrastive framework for neural text generation](#). *CoRR*, abs/2202.06417.
- Sainbayar Sukhbaatar, Edouard Grave, Piotr Bojanowski, and Armand Joulin. 2019. [Adaptive attention span in transformers](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 331–335. Association for Computational Linguistics.
- Zewei Sun, Shujian Huang, Hao-Ran Wei, Xin-yu Dai, and Jiajun Chen. 2020. [Generating diverse translation by manipulating multi-head attention](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8976–8983.
- Yi Tay, Dara Bahri, Liu Yang, Donald Metzler, and Da-Cheng Juan. 2020a. [Sparse sinkhorn attention](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 9438–9447. PMLR.

- Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. 2020b. [Efficient transformers: A survey](#). *CoRR*, abs/2009.06732.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Ashwin K. Vijayakumar, Michael Cogswell, Ramprasaath R. Selvaraju, Qing Sun, Stefan Lee, David J. Crandall, and Dhruv Batra. 2018. [Diverse beam search for improved description of complex scenes](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), the 30th Innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 7371–7379. AAAI Press.
- Ke Wang and Xiaojun Wan. 2018. Sentigan: Generating sentimental texts via mixture adversarial networks. In *International Joint Conference on Artificial Intelligence*.
- Shuohang Wang, Luowei Zhou, Zhe Gan, Yen-Chun Chen, Yuwei Fang, Siqi Sun, Yu Cheng, and Jingjing Liu. 2020. [Cluster-former: Clustering-based sparse transformer for long-range dependency encoding](#). *CoRR*, abs/2009.06097.
- Sean Welleck, Ilya Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2020. [Neural text generation with unlikelihood training](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Fangzhao Wu, Ying Qiao, Jiun-Hung Chen, Chuhan Wu, Tao Qi, Jianxun Lian, Danyang Liu, Xing Xie, Jianfeng Gao, Winnie Wu, and Ming Zhou. 2020. [MIND: A large-scale dataset for news recommendation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3597–3606, Online. Association for Computational Linguistics.
- Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tiejun Liu. 2020. [On layer normalization in the transformer architecture](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 10524–10533. PMLR.
- Jin Xu, Xiaojiang Liu, Jianhao Yan, Deng Cai, Huayang Li, and Jian Li. 2022a. Learning to break the loop: Analyzing and mitigating repetitions for neural text generation. *arXiv preprint arXiv:2206.02369*.
- Jin Xu, Xiaojiang Liu, Jianhao Yan, Deng Cai, Huayang Li, and Jian Li. 2022b. [Learning to break the loop: Analyzing and mitigating repetitions for neural text generation](#). *CoRR*, abs/2206.02369.
- Xiaoyuan Yi, Maosong Sun, Ruoyu Li, and Wenhao Li. 2018. Automatic poetry generation with mutual reinforcement learning. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3143–3153, Brussels, Belgium. Association for Computational Linguistics.
- Sangwon Yu, Jongyoon Song, Heeseung Kim, Seongmin Lee, Woo-Jong Ryu, and Sungroh Yoon. 2022a. [Rare tokens degenerate all tokens: Improving neural text generation via adaptive gradient gating for rare token embeddings](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 29–45, Dublin, Ireland. Association for Computational Linguistics.
- Sangwon Yu, Jongyoon Song, Heeseung Kim, Seongmin Lee, Woo-Jong Ryu, and Sungroh Yoon. 2022b. [Rare tokens degenerate all tokens: Improving neural text generation via adaptive gradient gating for rare token embeddings](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 29–45. Association for Computational Linguistics.
- Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontañón, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. 2020. [Big bird: Transformers for longer sequences](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Biao Zhang, Ivan Titov, and Rico Sennrich. 2021. [On sparsifying encoder outputs in sequence-to-sequence models](#). In *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, volume ACL/IJCNLP 2021 of *Findings of ACL*, pages 2888–2900. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. [Character-level Convolutional Networks for Text Classification](#). *arXiv:1509.01626 [cs]*.

Wangchunshu Zhou, Qifei Li, and Chenle Li. 2021. [Learning from perturbations: Diverse and informative dialogue generation with inverse adversarial training](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 694–703. Association for Computational Linguistics.

Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. 2018. [Txygen: A benchmarking platform for text generation models](#). In *International ACM SIGIR Conference on Research and Development in Information Retrieval*.

Appendix

A Experimental Details

A.1 Datasets

We conduct our experiments on four datasets, ROCStories, ParaSCI, Yelp, and MP (combing MIND & PENS). The statistics of the datasets are listed in Table 3, and the GPU hours for training and inference are in Table 4. All the datasets are in English.

ParaSCI (Dong et al., 2021a) A paraphrasing dataset collected from scientific literature, particularly ACL Anthology and arXiv. We use the ParaSCI-arXiv set since it contains more instances while following the split of their original paper.

Yelp (Zhang et al., 2015) A dataset for unconditional generation collected from the food review on the Yelp website. We also follow the split released publicly.

ROCStories (Mostafazadeh et al., 2016) A story corpus consists of 5-sentence short stories. We download it from the original website and randomly split it with the proportion in 4, following the previous papers (Guan et al., 2020). We use the first sentence as the condition to generate the following four sentences, which also agree with the settings in (Guan et al., 2020)

MP (MIND and PENS) (Wu et al., 2020; Ao et al., 2021) They are two datasets consisting of the main body and the headline of online news, providing a source for generating the headline of news given its content. Since they are both organized from MSN News, we merged them for a larger dataset and divided the training, validation, and test dataset with the proportions in 4. We truncate the news body to 1024 tokens since some are too long.

	Train	Inference
ROCStories	8h	7h
ParaSCI	5h	3h
MP	50h	5h
Yelp	15h	2h

Table 3: The training and inferring time consumed for each dataset. The inference time means the time consuming of the inference process of one checkpoint on the whole test dataset for one round with the corresponding decoding strategy (sampling for Yelp and beam search for others).

Dataset	#Train	#Dev	#Test
ROCStories	88345	4908	4908
ParaSCI	309834	3680	2549
MP	214964	5000	10000
Yelp	100000	10000	10000

Table 4: Dataset Statistics (Train/Dev/Test Split) for all the dataset we train on

A.2 Metrics

We use the following metrics to evaluate the quality of the generated text:

BLEU (Papineni et al., 2002) is a widely used metric in a variety of tasks on machine translation and language generation, calculating the n-gram overlap between the generated text and its corresponding reference. We use the implementation in the Huggingface Dataset Library(Lhoest et al., 2021). We reported BLEU-2 and BLEU-4 since they are most commonly used in language generation. In the unconditional generation task, we treat all the instances in the test set as references.

ROUGE (Lin and Hovy, 2002) is another metric based on the n-gram overlap between the generated and referenced sequences, mostly used in automatic summarization. We also use the version from

Huggingface Dataset Library(Lhoest et al., 2021). We measured ROUGE-2,3, L, and W on the generated texts.

BertScore(Zhang et al., 2020) is a model-based metric measuring the semantic similarity by the cosine similarity between generated sequence and reference using pre-trained BERT contextual representations. We use the *bert_score* library⁵ in our evaluation.

CND (Li et al., 2020b) is a distribution-level metric that approximates the divergence between the distribution of generated text and its corresponding reference in n-gram spaces. We use our self-implemented version of it.

MAUVE (Pillutla et al., 2021) is another distributional-level metric measuring the distributional divergence between the generated text and the reference text. We use its official implementation from GitHub⁶.

For diversity and novelty, we take the following three metrics:

JS (Wang and Wan, 2018) measures the diversity of generated text by calculating the Jaccard Similarity between the n-grams of each pair of generated sequences. We implement it by ourselves.

Dist (Li et al., 2016a) is a widely used metric for the novelty of the generated text, which calculates the proportion of distinct n-grams within the generated sentences. We also use the self-implemented version of it.

Self-BLEU (Zhu et al., 2018) measures the diversity of generated corpus by the average of pairwise BLEU score in the corpus. We use the version from the *fast_bleu* library⁷.

We reported the geometric mean of 1-4 grams for CND, JS, Dist, and Self-BLEU.

A.3 Implementation Details

We use GPT-2 base (Radford et al., 2019) as the backbone for ROCStories, ParaSCI, and Yelp, and a little smaller model for headline generation due to the limit of our computational power. The specific hyperparameter of the architecture of the model we demonstrate in 5. Models on the MP dataset is trained on NVIDIA GeForce GTX TITAN X, while models on the other three datasets is trained on NVIDIA GeForce RTX 3090.

For the datasets of conditional generation, we only compute loss on the continuation text but not on the conditional prefix. In addition, we train a BPE tokenizer for each dataset with the vocabulary size of 30000 to perform better tokenization.

We use beam search on the conditional generation dataset due to its stability and sampling for unconditional generation since beam search in this task tends to generate duplicated candidates. For the specific parameters in the unconditional generation, we use a hybrid strategy with top-p=0.9 (Holtzman et al., 2020) and top-k=50 (Fan et al., 2018) with temperature 1. We also perform a linear warmup on the \mathcal{L}_R term in CARE and the entropy term in CARE-A. Our methods and all the baselines are implemented by HuggingFace Transformers library (Wolf et al., 2020). All the models are trained from scratch.

	hidden_size	head_num	layer_num	Total_Parameters
MP	512	8	8	39.33M
Others	768	12	12	101.03M

Table 5: The detailed architecture of the model

⁵<https://pypi.org/project/bert-score/>

⁶<https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=&cad=rja&uact=8&ved=2ahUKewixy47LvMj4AhUeJ0QIHc-4CiUQFnoECAGQAQ&url=https%3A%2F%2Fgithub.com%2Fkrishnap25%2Fmauve&usg=AOvVaw2PEycTPemGnrAhnGagzJEF>

⁷<https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=&cad=rja&uact=8&ved=2ahUKewjc1NjWvcj4AhVTKEQIHxh-ADwQFnoECAKQAQ&url=https%3A%2F%2Fgithub.com%2FDanial-A1h%2Ffast-bleu&usg=AOvVaw0Dun5WnEw2LscEz7Gyhxeu>

A.4 Hyperparameter Tuning

We ran the model for different epochs reported in Table 6. We evaluate each epoch and report the best epoch due to the performance on the validation set. We search each model on each dataset for 30 hyperparameter trials. With two variants and four datasets, we ran our model about 240 times. We manually tune the models based on the overall performance of the model on the quality and diversity measurements of the validation set, mainly considering CND, BertScore, Dist, and Self-BLEU. The upper and lower bounds of the hyperparameters of our model are demonstrated in Table 7.

Dataset	Epoch
MP	20
ROC	25-28
ParaSCI	5
Yelp	25-30

Table 6: Running Epoch for different datasets

Hyperparameters	lower bound	upper bound
α	1.5	6.0
γ	1e-6	1e-2
δ	1e-5	1e-3
Learning Rate	5e-5	5e-3
Freezing Steps for \mathcal{L}_R	0	20000
Warmup Steps for \mathcal{L}_R	0	20000

Table 7: Bounds for the hyperparameter used

A.5 Human Evaluation

We randomly selected 100 prompts from the ROCStories dataset and let each model generate corresponding continuations. In the annotation sheet, we present the prompt and the continuations to the annotators in the format of 100 groups. Each group contains six continuations generated from the same prompt by the five models and the ground truth. The order of the model (or ground truth) within each group is randomly shuffled, and the model each continuation from is unknown to the annotator. The different orders in different groups could alleviate the bias of the annotators. We recruited four annotators with adequate English skills and asked them to score the fluency of the shuffled stories on a scale of 1-5 and calculate the average scoring for each model, as shown in Figure 6.

B Supplemental Materials for the Analysis of the *Trap of Medicority*

B.1 The Evidence that the Hidden State is nearer with the High-Frequency Word Embeddings

To prove the claim that the hidden states are closer to the high-frequency words, we calculate the pairwise cosine similarity between the hidden state when generation and the high-frequency and low-frequency word embeddings and demonstrate the result in Figure 9. The significant gap between these two lines proves our assumption that the high-frequency word embeddings are closer to the hidden states than the low-frequency ones.

B.2 The Addition Results on the Yelp Dataset

We conducted the same experiment in Section 2 on the Yelp Dataset. We correspondingly demonstrate the result of Figure 2, 3, 4, 5 and 9 in Figure 10, 11, 12, 13 and 14. The results all indicate the same result as in ROC stories, demonstrating the generalizability of our theory across different datasets and across the conditional and unconditional generation tasks.

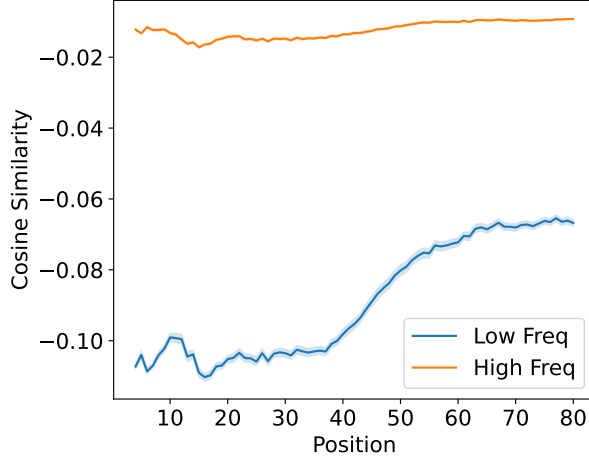


Figure 9: The pairwise cosine similarity between the hidden state when generation and the high-frequency and low-frequency word embeddings. The gap between two lines indicated that the high-frequency word embeddings are closer to the hidden states

C Derivation and Proof

C.1 The Demonstration for the Uniformly Positive Direction in Sub-claim 3

Gao et al. (2019) proves their Theorem 2 by separately proving the necessary and sufficient part. We will correspondingly modify them to prove the same conclusion in the uniformly positive direction.

In the necessary part, they begin by assuming the convex hull contains the origin. This assumption leads to the conclusion that there exists α_i and a vector \mathbf{v} satisfying $\sum_i \alpha_i \langle \mathbf{h}_i, \mathbf{v} \rangle = 0$, which contradicts the existence of the uniformly negative direction, *i.e.*, $\forall i, \langle \mathbf{h}_i, \mathbf{v} \rangle < 0$ for the hidden states \mathbf{h}_i . Meanwhile, this inducted form also contradicts the existence of the uniformly positive direction $\forall i, \langle \mathbf{h}_i, \mathbf{v} \rangle > 0$. Thus, the necessary part of the uniformly positive direction could be easily verified also by contradiction.

For the sufficient part, the authors illustrate the existence of a hyperplane containing the origin and apart from the convex hull. A normal direction of this hyperplane would be the uniformly negative direction. For modification, another normal direction opposite from the uniformly negative direction could also meet the requirement for the uniformly positive direction.

Combining the necessary and sufficient parts, we can know there is a uniformly positive direction for all hidden states. As the hidden states of frequent tokens account for the largest proportion of all hidden states (Pareto principle), such a uniform direction tends to be closer to hidden states corresponding to the frequent tokens, successfully proving sub-claim 3.

C.2 The Derivation of Sub-claim 4

For some specific word with index k in a sentence x with length m , assuming that this word only occurs on position m , that is, only $x_m = k$. Let $p_i = p(i = k | \mathbf{h}_i)$, then we have

$$\frac{\partial L_{NLL}}{\partial \mathbf{w}_k} = (1 - p_m) \mathbf{h}_m + \sum_{i=1}^{m-1} p_i \mathbf{h}_i. \quad (9)$$

Although we could yield the conclusion from this formula that the optimization direction of the word embedding agrees with the direction of hidden states (since $1 - p_m$ and p_i are all positive), the contributions of the hidden states are highly biased by p_i . In the early stage of training, $p_i \approx \frac{1}{V}$ for all i , thus the optimization direction will be mainly directed by \mathbf{h}_m . While in the late stage of training, $p_i \approx 0$ for $i \neq m$, and $p_{i \neq m} \ll p_m \approx \epsilon < 1$, the direction of \mathbf{h}_m will also dominate the optimization. We could barely tell that the optimization direction is the uniformly positive direction from the final layer. Thus we resort to the previous layer.

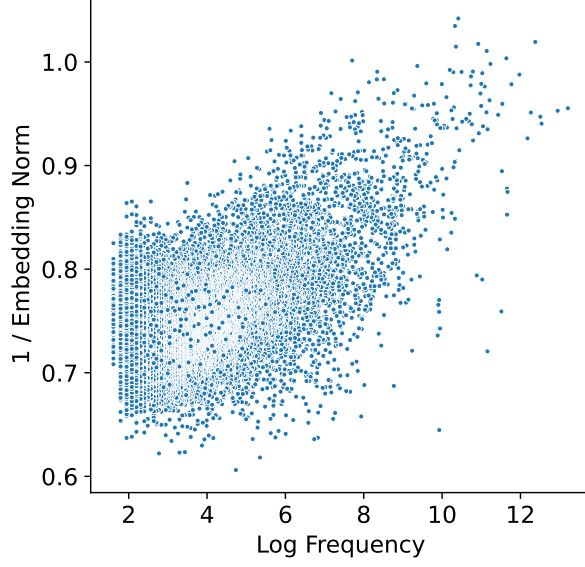


Figure 10: The correlation between reciprocal embedding norm and logarithmic word frequency on the Yelp Dataset, with the same setting in Figure 2. The corresponding R^2 is 0.854.

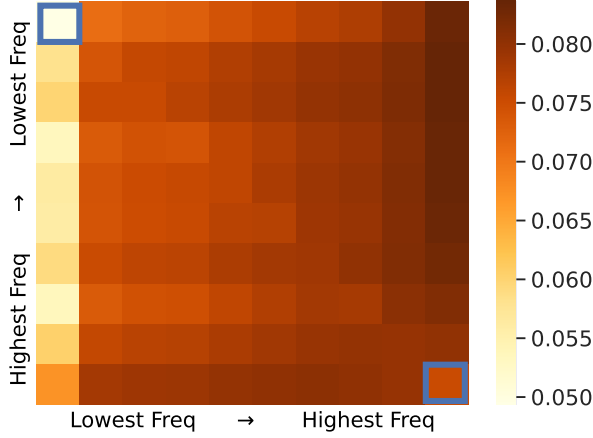


Figure 11: The distance of embeddings in/between different frequency intervals on the Yelp Dataset, with the same setting in Figure 3

For simplification, we ignore the nonlinear layer and let $\mathbf{h}_i = \sum_j a_{i,j} \hat{\mathbf{h}}_j$, where $\hat{\mathbf{h}}_j$ is the hidden state passed from the previous layer in Transformer, then we have:

$$\begin{aligned} \frac{\partial L_{NLL}}{\partial \mathbf{w}_k} &= \sum_{i=1}^{m-1} \left[a_{m,i}(1 - p_m) - \sum_{j=i}^{m-1} a_j^i p_j \right] \hat{\mathbf{h}}_i \\ &+ (1 - p_m) a_{m,m} \hat{\mathbf{h}}_m. \end{aligned} \quad (10)$$

We then consider the corner case when the attention distribution reaches the highest entropy, which

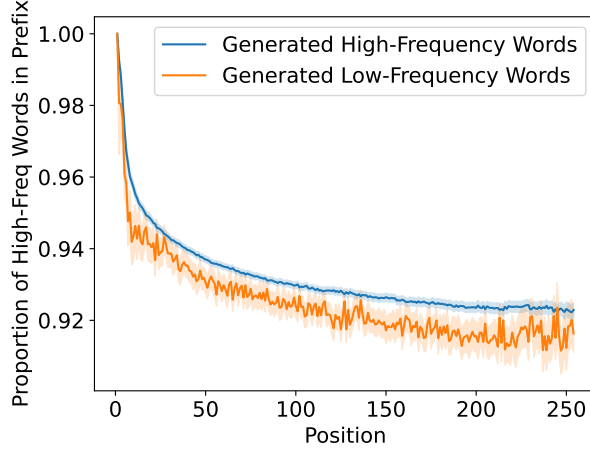


Figure 12: The distance of embeddings in/between different frequency intervals on the Yelp Dataset, with the same setting in Figure 4.

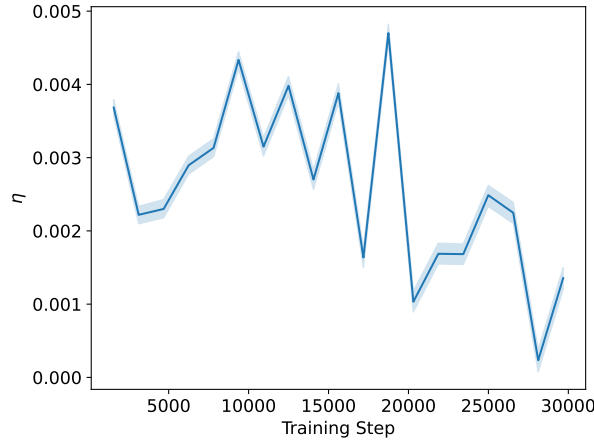


Figure 13: The value of η in Sub-claim 4 during training on the Yelp dataset, with the same setting in Figure 5.

leads to $a_{i,k} = \frac{1}{i}$. In this case, we have

$$\begin{aligned} \frac{\partial L_{NLL}}{\partial \mathbf{w}_k} &= \sum_{i=1}^{m-1} \left(\frac{1-p_m}{m} - \sum_{j=i}^{m-1} \frac{p_j}{j} \right) \hat{\mathbf{h}}_i \\ &+ \frac{1-p_m}{m} \hat{\mathbf{h}}_m \\ &= \sum_{i=1}^m \left(\frac{1}{m} - \sum_{j=i}^m \frac{p_j}{j} \right) \hat{\mathbf{h}}_i \end{aligned} \quad (11)$$

Thus, if $\frac{1}{m} - \sum_{j=i}^m \frac{p_j}{j} > 0$ always stands, the optimization direction of the embeddings would approach the mixture of the hidden states in the preceding context. If a uniformly positive direction of the hidden states exists, i.e. $\hat{\mathbf{h}}_i \mathbf{v}^T > 0$ stands for all the hidden states, then we have

$$\frac{\partial L_{NLL}}{\partial \mathbf{w}_k} \mathbf{v}^T = \sum_{i=1}^m \left(\frac{1}{m} - \sum_{j=i}^m \frac{p_j}{j} \right) \hat{\mathbf{h}}_i \mathbf{v}^T > 0. \quad (12)$$

Thus, the direction \mathbf{v} is also a uniformly positive direction for the optimization direction of the word embeddings. Moreover, the gradient contribution of $\hat{\mathbf{h}}_i$ is more scattered among the hidden states, and the tendency will be advanced when considering more previous layers. Therefore, we could conclude that the optimization will approach the uniformly positive direction under this condition.

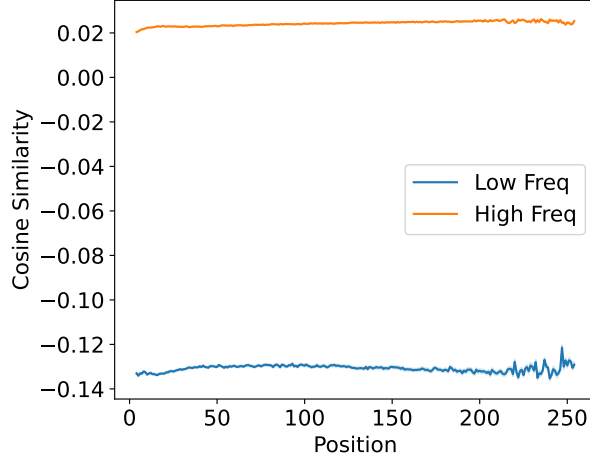


Figure 14: The pairwise cosine similarity between the hidden state when generation and the high-frequency and low-frequency word embeddings in the Yelp dataset, with the same setting in Figure 9.

C.3 The Proof for Theorem 1

We start our proof by the definition of the Rényi entropy

$$\begin{aligned}
 H_\alpha(\mathbf{a}_t) &= \frac{1}{1-\alpha} \log\left(\sum_i a_{t,i}^\alpha\right) \\
 &= \frac{1}{1-\alpha} \log\left(\sum_i \frac{e^{\alpha \tilde{a}_{t,i}}}{Z^\alpha}\right) \\
 &= \frac{1}{1-\alpha} \left[\log\left(\sum_i e^{\alpha \tilde{a}_{t,i}}\right) - \alpha \log Z \right]
 \end{aligned} \tag{13}$$

Then we let $\alpha > 1$, then

$$H_\alpha(\mathbf{a}_t) = \frac{1}{\alpha-1} [\alpha \log Z - \log(\sum_i e^{\alpha \tilde{a}_{t,i}})] \tag{14}$$

According to the Jensen Inequality,

$$\begin{aligned}
 \log\left(\sum_i e^{\alpha \tilde{a}_{t,i}}\right) &= \log\left(\sum_i \frac{e^{\alpha \tilde{a}_{t,i}}}{t}\right) + \log t \\
 &\geq \frac{\sum_i \log e^{\alpha \tilde{a}_{t,i}}}{t} + \log t \\
 &= \frac{\alpha \sum_i \tilde{a}_{t,i}}{t} + \log t
 \end{aligned} \tag{15}$$

$$\begin{aligned}
 \therefore H_\alpha(\mathbf{a}) &\leq \frac{1}{\alpha-1} \left[\alpha \log Z - \frac{\alpha}{t} \sum \tilde{a}_{t,i} - \log t \right] \\
 &\leq \frac{1}{\alpha-1} \left[\alpha \log Z - \frac{\alpha}{t} \sum \tilde{a}_{t,i} \right]
 \end{aligned} \tag{16}$$

While

$$\begin{aligned}
 \log Z &= \log \sum e^{\tilde{a}_{t,i}} \\
 &\leq \max\{\tilde{a}_{t,1}, \dots, \tilde{a}_{t,n}\} + \log t \\
 &\leq \max\{|\tilde{a}_{t,1}|, \dots, |\tilde{a}_{t,n}|\} + \log t \\
 &= \|\tilde{\mathbf{a}}_t\|_\infty + \log t
 \end{aligned} \tag{17}$$

$$\begin{aligned}
\therefore H_\alpha(\mathbf{a}_t) &\leq \frac{1}{\alpha-1} \left[\alpha \|\tilde{\mathbf{a}}_t\|_\infty + \alpha \log t - \frac{\alpha}{t} \sum_i \tilde{a}_{t,i} - \log t \right] \\
&\leq \frac{\alpha}{\alpha-1} \left[\|\tilde{\mathbf{a}}_t\|_\infty - \frac{1}{t} \sum_i \tilde{a}_{t,i} \right] + \log t \\
&\leq \frac{\alpha}{\alpha-1} \left[\|\tilde{\mathbf{a}}_t\|_1 - \frac{1}{t} \sum_i \tilde{a}_{t,i} \right] + \log t
\end{aligned} \tag{18}$$

Then we consider the term

$$\begin{aligned}
\|\tilde{\mathbf{a}}_t\|_1 - \frac{1}{t} \sum_i \tilde{a}_{t,i} &= \sum_i \left[|\tilde{a}_{t,i}| - \frac{1}{t} \tilde{a}_{t,i} \right] \\
&\leq \sum_i \left[\left(1 + \frac{1}{t}\right) |\tilde{a}_{t,i}| \right]
\end{aligned} \tag{19}$$

$$\begin{aligned}
\therefore H_\alpha(\mathbf{a}_t) &\leq \frac{\alpha \left(1 + \frac{1}{t}\right)}{\alpha-1} \sum_i |\tilde{a}_{t,i}| + \log t \\
&= \frac{\alpha \left(1 + \frac{1}{t}\right)}{\alpha-1} \|\tilde{\mathbf{a}}_t\|_1 + \log t \\
&= \frac{\alpha(t+1)}{t(\alpha-1)} \|\tilde{\mathbf{a}}_t\|_1 + \log t \\
&\leq \frac{\alpha(t+1)}{t(\alpha-1)} \|\hat{\mathbf{a}}_t\|_1 + \log t + \frac{\alpha(t+1)}{\alpha-1} |C|
\end{aligned} \tag{20}$$

where C is a large negative constant used in attention dropout.

Discarding all the constants, minimizing the upper bound of $H_\alpha(\mathbf{a}_t)$ above is equal to minimize

$$\frac{\alpha(t+1)}{t(\alpha-1)} \|\hat{\mathbf{a}}_t\|_1 \tag{21}$$

C.4 The Proof for Theorem 2

We use the Bayes theorem to formalize the posterior distribution as follows:

$$\begin{aligned}
&KL[q_\theta(\tilde{\mathbf{a}}_t) \| p(\tilde{\mathbf{a}}_t | x_t, y_t, c)] \\
&= \int q_\theta(\tilde{\mathbf{a}}_t) \log \frac{q_\theta(\tilde{\mathbf{a}}_t)}{p(\tilde{\mathbf{a}}_t | x_t, y_t, c)} d\tilde{\mathbf{a}}_t \\
&= \int q_\theta(\tilde{\mathbf{a}}_t) \log \frac{q_\theta(\tilde{\mathbf{a}}_t)}{p(\tilde{\mathbf{a}}_t | x_t, y_t, c)} \cdot \frac{p(\tilde{\mathbf{a}}_t)}{p(\tilde{\mathbf{a}}_t)} d\tilde{\mathbf{a}}_t \\
&= KL[q_\theta(\tilde{\mathbf{a}}_t) \| p(\tilde{\mathbf{a}}_t)] + \int q_\theta(\tilde{\mathbf{a}}_t) \log \frac{p(\tilde{\mathbf{a}}_t)}{p(\tilde{\mathbf{a}}_t | x_t, y_t, c)} d\tilde{\mathbf{a}}_t \\
&= KL[q_\theta(\tilde{\mathbf{a}}_t) \| p(\tilde{\mathbf{a}}_t)] - \int q_\theta(\tilde{\mathbf{a}}_t) \log \frac{p(\tilde{\mathbf{a}}_t | x_t, y_t, c)}{p(\tilde{\mathbf{a}}_t)} d\tilde{\mathbf{a}}_t
\end{aligned} \tag{22}$$

$$\begin{aligned}
\because p(\tilde{\mathbf{a}}_t | x_t, y_t, c) &= \frac{p(\tilde{\mathbf{a}}_t) p(x_t, y_t, c | \tilde{\mathbf{a}}_t)}{p(x_t, y_t, c)} \\
\therefore \frac{p(\tilde{\mathbf{a}}_t | x_t, y_t, c)}{p(\tilde{\mathbf{a}}_t)} &= \frac{p(x_t, y_t, c | \tilde{\mathbf{a}}_t)}{p(x_t, y_t, c)} \\
&= \frac{p(y_t | x_t, c, \tilde{\mathbf{a}}_t) p(x_t, c | \tilde{\mathbf{a}}_t)}{p(y_t | x_t, c) p(x_t, c)} \\
&= \frac{p(y_t | x_t, c, \tilde{\mathbf{a}}_t)}{p(y_t | x_t, c)}
\end{aligned} \tag{23}$$

$$\begin{aligned}
& \therefore KL [q_\theta(\tilde{\mathbf{a}}_t) || p(\tilde{\mathbf{a}}_t | x_t, y_t, c)] \\
& = KL [q_\theta(\tilde{\mathbf{a}}_t) || p(\tilde{\mathbf{a}}_t)] - \int q_\theta(\tilde{\mathbf{a}}_t) \log \frac{p(y_t | x_t, c, \tilde{\mathbf{a}}_t)}{p(y_t | x_t, c)} d\tilde{\mathbf{a}}_t \\
& = KL [q_\theta(\tilde{\mathbf{a}}_t) || p(\tilde{\mathbf{a}}_t)] - E_{q_\theta(\tilde{\mathbf{a}}_t)} [\log p(y_t | x_t, c, \tilde{\mathbf{a}}_t)] + \log p(y_t | x_t, c)
\end{aligned} \tag{24}$$

Thus, the process to minimize the KL Divergence $KL [q_\theta(\tilde{\mathbf{a}}_t) || p(\tilde{\mathbf{a}}_t | x_t, y_t, c)]$ is equivalent to maximize

$$\begin{aligned}
& - KL [q_\theta(\tilde{\mathbf{a}}_t) || p(\tilde{\mathbf{a}}_t | x_t, y_t, c)] \\
& = E_{q_\theta(\tilde{\mathbf{a}}_t)} [\log p(y_t | x_t, c, \tilde{\mathbf{a}}_t)] - KL [q_\theta(\tilde{\mathbf{a}}_t) || p(\tilde{\mathbf{a}}_t)] - \log p(y_t | x_t, c) \\
& \geq E_{q_\theta(\tilde{\mathbf{a}}_t)} [\log p(y_t | x_t, c, \tilde{\mathbf{a}}_t)] - KL [q_\theta(\tilde{\mathbf{a}}_t) || p(\tilde{\mathbf{a}}_t)]
\end{aligned} \tag{25}$$

We denote

$$q_\theta(\tilde{\mathbf{a}}_t) = p\delta(\tilde{\mathbf{a}}_t - C) + (1 - p)\delta(\tilde{\mathbf{a}}_t - \hat{\mathbf{a}}_t) \tag{26}$$

where C is a large negative number, and p is the probability of the attention dropout. We define $p(\tilde{\mathbf{a}}_t) \propto 1 + e^{-\gamma\|\tilde{\mathbf{a}}_t\|}$, then we have

$$\begin{aligned}
KL [q_\theta(\tilde{\mathbf{a}}_t) || p(\tilde{\mathbf{a}}_t)] & = \int q_\theta(\tilde{\mathbf{a}}_t) \log \frac{q_\theta(\tilde{\mathbf{a}}_t)}{p(\tilde{\mathbf{a}}_t)} d\tilde{\mathbf{a}}_t \\
& = \int q_\theta(\tilde{\mathbf{a}}_t) \log \frac{q_\theta(\tilde{\mathbf{a}}_t)}{1 + e^{-\gamma\|\tilde{\mathbf{a}}_t\|}} d\tilde{\mathbf{a}}_t + \log Z \\
& = p \log p + (1 - p) \log \frac{1 - p}{1 + e^{-\gamma\|\hat{\mathbf{a}}_t\|}} + \log Z \\
& = p \log p + (1 - p) \log(1 - p) + (1 - p) \log(1 + e^{-\gamma\|\hat{\mathbf{a}}_t\|}) + \log Z \\
& = -H(p) + (1 - p) \log(1 + e^{-\gamma\|\hat{\mathbf{a}}_t\|}) + \log Z
\end{aligned} \tag{27}$$

According to the Jensen Inequality, we have

$$\begin{aligned}
\frac{e^x + e^{-x}}{2} & \geq e^{\frac{x}{2}} \\
\therefore 1 + e^x & \geq 2e^{\frac{x}{2}} \\
\therefore \log(1 + e^x) & \geq \log 2 + \frac{x}{2}
\end{aligned} \tag{28}$$

$$\begin{aligned}
\therefore KL [q_\theta(\tilde{\mathbf{a}}_t) || p(\tilde{\mathbf{a}}_t)] & \leq -H(p) - (1 - p) \left[\log 2 + \frac{-\gamma\|\hat{\mathbf{a}}_t\|}{2} \right] + \log Z \\
& = -H(p) + \frac{\gamma(1 - p)}{2} \|\hat{\mathbf{a}}_t\| - (1 - p) \log 2 + \log Z
\end{aligned} \tag{29}$$

Discard all the constants, and we have

$$-H(p) + \frac{\gamma(1 - p)}{2} \|\hat{\mathbf{a}}_t\| \tag{30}$$

D Additional Results

D.1 Result with CARE-O

In Section 3, we mentioned a variant of our method only with the \mathcal{L}_R term remaining the original dropout pattern. In this section, we call this variant CARE-O, and examine its performance on ROC, Yelp, and ParaSCI. The results are listed in Table 8 and Table 9. As we can see, although it could achieve the best in some metrics like CND in ParaSCI, its overall quality is worse than GPT-2 only with a small margin on diversity improvement on ROC and ParaSCI. It also performs very poorly on Yelp. This experimental result agrees with our theoretical induction of the necessity of dropout modification.

Model	Quality								Diversity		
	R-2↑	R-3↑	R-L↑	R-W↑	B-2↑	B-4↑	BS↑	CND↓	Dist↑	JS↓	SB↓
Dataset: ParaSCI											
GPT-2	41.48	32.65	54.97	35.28	41.64	<u>26.58</u>	90.81	1.688	60.34	0.0726	17.52
Pattern	41.46	32.58	55.06	35.34	41.31	26.43	<u>90.89</u>	1.652	60.65	0.0683	16.86
Entmax	38.44	29.78	52.28	33.50	38.06	23.60	90.35	1.769	59.30	0.0707	18.10
l_0 -Drop	37.26	28.64	50.99	32.53	37.96	23.47	90.05	1.700	60.94	<u>0.0591</u>	<u>16.18</u>
\mathcal{L}_A -Tuning	40.71	32.05	54.20	34.89	39.95	25.48	90.70	1.632	60.67	0.0622	16.28
CARE-O	41.06	32.59	54.16	35.09	38.30	24.28	90.61	1.626	59.21	0.0685	17.28
CARE	42.49	33.74	<u>55.69</u>	35.94	41.59	26.75	90.80	<u>1.631</u>	61.04	0.0566	16.05
CARE-A	<u>41.95</u>	<u>33.01</u>	55.79	<u>35.76</u>	<u>41.63</u>	26.56	91.04	1.637	60.60	0.0596	17.04
Dataset: ROCStories											
GPT-2	6.700	1.321	23.17	<u>11.19</u>	5.817	0.266	83.53	8.856	25.72	0.6625	62.36
Pattern	6.690	1.301	23.33	11.09	6.207	0.273	<u>83.64</u>	9.006	27.20	0.6280	61.16
Entmax	6.011	1.011	22.42	10.80	5.374	0.150	82.91	9.521	25.03	0.7317	62.57
l_0 -Drop	5.850	1.065	22.11	10.77	4.896	0.177	83.17	8.489	24.82	0.6207	62.72
\mathcal{L}_A -Tuning	6.367	1.214	22.94	10.96	5.809	0.226	83.37	8.904	<u>27.24</u>	<u>0.5905</u>	61.06
CARE-O	5.984	1.068	22.52	10.78	5.401	0.199	83.41	<u>8.603</u>	26.78	0.5915	<u>61.05</u>
CARE	<u>6.905</u>	1.399	<u>23.54</u>	11.08	6.829	0.326	83.74	9.735	27.77	0.6328	60.41
CARE-A	7.002	<u>1.332</u>	23.80	11.32	<u>6.489</u>	<u>0.274</u>	83.63	8.742	25.89	0.5708	61.51

Table 8: The evaluation result of conditional generation with CARE-O

Model	Quality				Diversity		
	B-2↑	B-4↑	CND↓	MV↑	Dist↑	JS↓	SB↓
GPT-2	76.68	36.46	<u>0.886</u>	<u>0.931</u>	27.62	0.187	47.76
Pattern	76.96	<u>38.09</u>	0.924	0.859	27.34	0.203	49.31
Entmax	77.03	32.18	0.951	0.920	23.99	0.222	49.18
l_0 -Drop	76.71	34.92	0.866	0.919	26.33	0.194	47.38
\mathcal{L}_A -Tuning	<u>77.35</u>	38.66	0.910	0.835	26.26	0.211	50.82
CARE-O	76.31	26.87	1.262	0.5128	17.10	<u>0.224</u>	50.53
CARE	73.84	33.61	1.033	0.863	30.84	<u>0.178</u>	43.60
CARE-A	77.53	35.69	0.933	0.974	<u>28.44</u>	0.169	<u>44.65</u>

Table 9: Evaluation results for unconditional generation on the Yelp dataset with CARE-O. MV stands for MAUVE.

D.2 Result on MP

We also examine our model on MP dataset. The result is in 10 here due to the space limit. As we can see, our model achieves the best diversity and novelty with a little sacrifice of the generated quality, still demonstrating its ability to enhance the novelty and the diversity of the generated text.

D.3 Result on the repetition

Although our model focus on improving inter-instance diversity instead of alleviating the intra-sentence repetition problem, we still add the result of Rep to examine the ability of our model to relieve the repetition. First used in (Welleck et al., 2020), Rep measures the percentage of repetitive generated n-grams. We implement this by ourselves and present the result of the geometric mean of Rep1-4 in ROCStories Dataset in Table 11.

As demonstrated, the method with sparse attention can achieve a better score of Rep, proving that sparse attention can enhance inter-instance diversity and soothe intra-sentence degeneration. Among them, our CARE and Pattern achieve significantly superior scores compared to others, demonstrating the CARE’s another efficacy in the reduction of repetition.

Model	Quality								Diversity		
	R-2 \uparrow	R-3 \uparrow	R-L \uparrow	R-W \uparrow	B-2 \uparrow	B-4 \uparrow	BS \uparrow	CND \downarrow	Dist \uparrow	JS \downarrow	SB \downarrow
GPT-2	<u>20.18</u>	<u>12.76</u>	<u>34.14</u>	<u>24.47</u>	<u>18.65</u>	<u>9.817</u>	<u>88.17</u>	0.468	62.29	0.0159	17.17
Pattern	20.34	12.82	34.27	24.59	18.77	9.850	88.23	0.470	62.26	0.0157	17.24
l_0 -Drop	19.65	12.00	33.78	24.14	18.21	8.838	88.14	0.480	62.38	0.0163	17.78
CARE-O	19.84	12.60	33.59	24.14	18.29	9.641	88.04	0.383	62.77	<u>0.0138</u>	16.98
CARE	19.49	12.31	33.24	23.87	17.98	9.625	88.02	<u>0.403</u>	<u>62.76</u>	0.0140	<u>16.90</u>
CARE-A	18.84	11.73	32.61	23.40	17.14	8.703	87.79	0.407	62.73	0.0129	15.88

Table 10: The evaluation result on headline generation.

Model	Rep \downarrow
GPT-2	5.598
Pattern	4.571
Entmax	5.170
l_0 -Drop	5.479
\mathcal{L}_A -Tuning	6.005
CARE	<u>4.742</u>
CARE-A	5.479

Table 11: Evaluation results of Rep on ROCStories Dataset.

E Python Implementation of CARE

We list the code of CARE as follows. This code is modified from the Huggingface Transformers Library (Wolf et al., 2020). Apart from the necessary function name and class name for location, the newly added code contains 19 added or modified lines highlighted in blue.

```
1 class GPT2Attention(nn.Module):
2     def __init__(self, config, is_cross_attention=False):
3         ....
4         self.attn_drop_ratio = config.attn_pdrop
5         ....
6     def _attn(self, query, key, value, attention_mask=None, head_mask=None):
7         ret_attn_weights = attn_weights
8         if self.training:
9             drop_mask = torch.zeros_like(attn_weights).bernoulli_(self.attn_drop_ratio)
10            attn_weights = nn.Softmax(dim=-1)(attn_weights + drop_mask * (-1e4))
11        else:
12            attn_weights = nn.Softmax(dim=-1)(attn_weights)
13        ....
14        return attn_output, ret_attn_weights
15 class CARETrainer(Seq2seqTrainer)
16     ....
17     def compute_loss(self, model, inputs, return_outputs=False):
18         ....
19         k_scale = torch.arange(1, outputs.attentions.size(-1) + 1,
20            ↪ device=inputs['attention_mask'].device, dtype=torch.float)
21         k_scale = k_scale.view(1, 1, -1).repeat(outputs.attentions.size(0),
22            ↪ outputs.attentions.size(1), 1)
23         len_vec = inputs['attention_mask'].sum(-1, keepdim=True).repeat(1, attn_weights[0].size(1))
24         scale = (1.0 + 1.0 / k_scale) * (1 - self.extra_args.attn_drop) * (1.0 + 1.0 / alpha)
25         attn_l1_loss = []
26         for layer_id in range(0, len(attn_weights)):
27             abs_attns = torch.abs(torch.tril(attn_weights[layer_id], 0)).sum(dim=-1)
28             attn_l1 = torch.div(torch.mul(abs_attns, scale).sum(-1), len_vec)
29             attn_l1_loss.append(attn_l1.mean().view(1))
30         loss += beta * torch.cat(attn_l1_loss).mean()
31     ....
```
