

# TEncDM: Understanding the Properties of the Diffusion Model in the Space of Language Model Encodings

Alexander Shabalin<sup>1,5\*</sup>, Viacheslav Meshchaninov<sup>1\*</sup>, Egor Chimbulatov<sup>1</sup>, Vladislav Lapikov<sup>1</sup>, Roman Kim<sup>1</sup>, Grigory Bartosh<sup>2</sup>, Dmitry Molchanov<sup>3</sup>, Sergey Markov<sup>4</sup>, Dmitry Vetrov<sup>5</sup>

<sup>1</sup>HSE University

<sup>2</sup>University of Amsterdam

<sup>3</sup>Independent Researcher

<sup>4</sup>SberDevices

<sup>5</sup>Constructor University

{amshabalin, vmeshchaninov, echimbulatov, vlapikov, rkim}@hse.ru, dmolch111@gmail.com, g.bartosh@uva.nl, sergei.markoff@gmail.com, dvetrov@constructor.university

## Abstract

This paper presents the *Text Encoding Diffusion Model* (TEncDM), a novel approach to diffusion modeling that operates in the space of pre-trained language model encodings. In contrast to traditionally used embeddings, encodings integrate contextual information. In our approach, we also employ a transformer-based decoder, specifically designed to incorporate context in the token prediction process. We conduct a comprehensive examination of the influence of the encoder, decoder, noise scheduler, and self-conditioning on zero-shot generation. Furthermore, we compare TEncDM with previous approaches on three conditional text generation tasks: QQP, XSum, and Wiki-Auto. The results show that TEncDM exhibits superior performance compared to existing non-autoregressive diffusion models.

**Code** — <https://github.com/M0RJIQUE/tencdm>

**Extended version** — <https://arxiv.org/abs/2402.19097>

## 1 Introduction

Autoregressive (AR) large language models such as GPT-4 (OpenAI 2023) or Llama 3 (Dubey et al. 2024) are the current gold standard in the text generation problem. They are capable of creating high-quality and coherent texts that are practically indistinguishable from the human ones. However, the disadvantage of this approach is the inability of the model to correct its own mistakes made during left-to-right generation. If a mistake occurs, it may spoil the subsequent text. In addition, the autoregressive method of token generation slows down the inference process as it requires performing a single model evaluation for each new token.

Diffusion modeling is currently the state-of-the-art approach for data generation in image (Rombach et al. 2022; Podell et al. 2023), audio (Evans et al. 2024) and video (Blattmann et al. 2023) domains. Researchers now attempt to adapt it also for text generation (Li et al. 2022; Gong et al. 2023; Karimi Mahabadi et al. 2024). Diffusion models are a

class of probabilistic generative models that are able to iteratively transfer noise to a representative sample of data. While some of the proposed text diffusion models are autoregressive (Lovelace et al. 2023; Zhang et al. 2023), the majority of them are not and, by design, they have several advantages over AR language models. First, being non-autoregressive (NAR) models, they generate all the tokens simultaneously and can adjust any part of the sequence during the generation process. They also can be faster than AR models because the number of neural function evaluations for diffusion models depends on the number of denoising iterations rather than the length of the sequence. And given the possibility of distillation of diffusion models (Meng et al. 2023), the number of iterations can be greatly reduced.

To date, a number of text diffusion models have been proposed, each based on substantially new ideas with little overlap with other methods. Some works replace Gaussian noise with categorical noise (Hoogetboom et al. 2021; Austin et al. 2021), exploiting the discreteness of the text domain. Others train continuous diffusion on token embeddings (Li et al. 2022; Lin et al. 2023; Gong et al. 2023), on text latent representations reduced in size (Lovelace et al. 2023; Zhang et al. 2023) or on the probability simplex (Karimi Mahabadi et al. 2024; Han, Kumar, and Tsvetkov 2023). There are also differences in the way diffusion outputs are decoded back into text. Diffusion models trained on embeddings round their predictions to the nearest embeddings, while those that utilize small latent spaces decode the predictions with an AR model. This suggests that the scientific community has not yet found the most robust design of a diffusion model for text.

In this paper, we attempt to better understand the specifics of continuous text diffusion models trained in the latent space of token embeddings and identify best practices for their development. We argue that the use of raw embeddings as latents is suboptimal, and that this approach can be enhanced by first extracting context information from embeddings with a language model encoder and training a diffusion model in this latent space. We also investigate several diffusion components in detail: text decoding meth-

\*First two authors contributed equally.

ods, diffusion model architecture, noise schedule, and self-conditioning (Chen, Zhang, and Hinton 2023). As a result, we combine all our findings in a method called *Text Encoding Diffusion Model (TEncDM)*.

We compare our approach with other non-autoregressive diffusion models trained in the space of embeddings or encodings on three conditional text generation problems – **paraphrasing**, **summarization**, and **text simplification** – and show its superiority over other methods. The main contributions of this work are as follows:

- We propose a new text diffusion framework **TEncDM**, which trains the diffusion model in the latent space constructed by the outputs of pre-trained Transformer-based encoder.
- We evaluate the importance of the text decoder and conclude that its robustness to inaccuracies in the generated latents directly affects the generation quality. We propose a Transformer-based decoder and its training method that boosts the model performance.
- We analyse in detail the effect of self-conditioning and noise schedules on the denoising process and show their effect on the model quality.

## 2 Problem Statement and Background

**Text generation problem.** In the field of natural language processing, unconditional text generation is a task of sampling  $y$  from the unknown distribution  $p(y)$ , where  $y = [y_1, \dots, y_n]$  is a sequence of tokens with variable length  $n$ . In conditional text generation the distribution of texts changes to  $p(y|x)$ , where  $x$  is a condition variable. The goal is to generate a text, that satisfies this condition.

**Gaussian diffusion models.** Gaussian diffusion models (Song et al. 2021) learn to sample data from an unknown distribution by gradually denoising random Gaussian noise. The training procedure is defined through a forward diffusion process that satisfies  $q(z_t|z_0) = \mathcal{N}(\sqrt{\alpha_t}z_0, (1 - \alpha_t)\mathbf{I})$ , where  $\alpha_t \in [0, 1]$  is a predefined noise schedule,  $t \in [0, 1]$  and  $\alpha_t > \alpha_{t+\Delta t}$ . The denoising network (parameterized by  $\theta$ ) is trained to reconstruct the original latent  $z_0$  given the noisy latent  $z_t$ , as expressed in equation 1.

$$\mathcal{L}(\theta) = \mathbb{E}_{\varepsilon \sim \mathcal{N}(0, \mathbf{I}), t \sim U[0;1]} [\|z_0 - \hat{z}_\theta(z_t, t)\|^2] \quad (1)$$

Sampling procedure starts from a pure Gaussian noise  $z_T \sim \mathcal{N}(0, \mathbf{I})$  and utilizes the denoising network to iteratively generate latents  $z_{t_{T-1}}, \dots, z_{t_1}$ , where  $1 = t_T > \dots > t_1 = 0$ .

**Diffusion models for text generation.** The primary feature of the text domain is the discreteness of its samples. In order to train a diffusion model on them, they must first be translated into continuous space. Consequently, alongside the denoising model, the diffusion framework incorporates an *encoder* that maps tokens into the continuous latents and a *decoder* that performs the reverse operation, converting the generated latents into text.

## 3 Related Work

**Embedding-based diffusion models.** The majority of proposed text diffusion models use embeddings of tokens

to construct the continuous latent space (Li et al. 2022; Lin et al. 2023; Strudel et al. 2022; Gong et al. 2023; Wu et al. 2023). At the inference stage, to convert the latent predictions into text, they map each latent vector to a token corresponding to the nearest embedding.

**Self-Conditioning.** Self-conditioning is a technique that significantly increases the performance of the text diffusion model (Chen, Zhang, and Hinton 2023; Strudel et al. 2022; Lovelace et al. 2023). Usually the model is conditioned only on the latent variable  $z_t$  and the current timestep  $t$  as  $\hat{z}_0^t = \hat{z}_\theta(z_t, t)$ . Self-conditioning proposes to also condition the model on the estimation of data sample from the previous timestep during generation in order to improve the prediction at the current timestep,  $\hat{z}_0^t = \hat{z}_\theta(z_t, t, \hat{z}_0^{t-1})$ .

Although widely used, no analysis has been conducted to determine why this method is effective or how the generation process is altered by its application.

**Noise scheduler.** Noise scheduler is a key component of a diffusion model that controls the amount of noise added on each timestep. Previous research (Li et al. 2022; Gao et al. 2024; Ye et al. 2024) has highlighted that the standard noise schedulers used for image diffusion models are unsuitable for the textual domain. Due to the discrete nature of the texts, it is unlikely that an addition of a small amount of noise to a latent will change its nearest text in the latent space. Therefore, to increase the difficulty of the denoising task for the model, the mentioned works recommend adding more noise on iterations that are close to 0.

## 4 Understanding Text Diffusion

In this section, we present our findings on the components of the diffusion model, discuss their weaknesses and propose ways to enhance them.

**Encodings are better than embeddings.** Most diffusion models utilize token embeddings to map text into a continuous latent space. However, this approach is not optimal because the embeddings do not convey contextual information. This requires the diffusion model to independently search for it to retrieve ambiguous tokens. To simplify the task, instead of embeddings, we can use the final layer outputs of a pre-trained language model (e.g. BERT). They contain contextual information and, thus, should be more suitable for training the diffusion model. We refer to these outputs as *encodings*.

Experimental results confirming our intuition are presented in Section 7.2. It is worth noting that the use of encodings does not slow down the generation process, as we need to compute them only during the training. To improve the quality even further, it is possible to fine-tune the encoder, but we choose not to in order to avoid overcomplicating the approach. Investigation into fine-tuning is left for the future work.

**Decoder is important.** The purpose of the decoder in the diffusion model is to map the generated latents into text. Approaches that train diffusion in the space of token embeddings decode latents by rounding them to the nearest

embeddings and selecting a corresponding token. However, the diffusion model may produce inaccurate latent samples due to accumulation of errors during the denoising process. Such inaccuracy might significantly spoil the text quality, so it would be wise to train a decoder that could improve it.

In the Section 7.2, we compare different decoder designs and conclude that an advanced context-dependent decoder indeed improves the generation quality.

**Self-conditioning affects denoising dynamics.** Self-conditioning improves sampling quality by conditioning the model on its previous prediction. However, the mechanics of self-conditioning are not fully understood yet. Our research demonstrates that the addition of self-conditioning increases the model’s prediction confidence at each denoising timestep, resulting in a reduction in the required number of generation steps. Furthermore, the sample quality diminishes as the number of steps increases. We believe that a reason for this behaviour lies in a mismatch between the latents used at the training stage and those at the generation stage. We provide the evidence supporting our conclusions in Section 7.2, along with a comprehensive analysis of the model’s behaviour with and without self-conditioning.

**Diffusion needs even more noise.** Following the recommendations of previous works (Li et al. 2022; Wu et al. 2023; Ye et al. 2024), we used *sqrt* noise scheduler that increases the amount of noise added to the diffusion model inputs during training beyond the amount of typically used *cosine* noise scheduler (Han, Kumar, and Tsvetkov 2023; Lovelace et al. 2023; Strudel et al. 2022; Zhang et al. 2023). However, our experiments led us to conclusion that encoding-based diffusion model requires even more noise for successful training. We hypothesize that this is due to the presence of contextual information in the encodings, which simplifies the denoising task.

In Section 7.2 of this study, we demonstrate that both commonly used *cosine* and *sqrt* noise schedules do not introduce a significant level of noise to the latent variables over a wide range of timesteps. As a result, the denoising task becomes too simple for the model, leading to a reduction in the effectiveness of the training signal.

## 5 Methodology

The design of **TEncDM** is depicted on Figure 1. It consists of three parts – diffusion encoder  $E_{diff}$ , diffusion model  $\hat{z}_\theta$  and decoder  $D$ . For the conditional generation, we also add condition encoder  $E_{cond}$ , which encodes an input text. Its output is provided to the diffusion model and decoder through cross-attention.

This section exclusively focuses on the topic of unconditional text generation. The details of the conditional model can be found in Section 5.4.

### 5.1 Diffusion encoder, $E_{diff}$

We use pre-trained Transformer-based (Vaswani et al. 2017b) language model  $E_{diff}$ , which we call *diffusion encoder*, to encode text  $y$  into the latent variable  $z$ . Encoding of text does not change the length of the sequence. In order to align all texts in length, we add paddings to the end

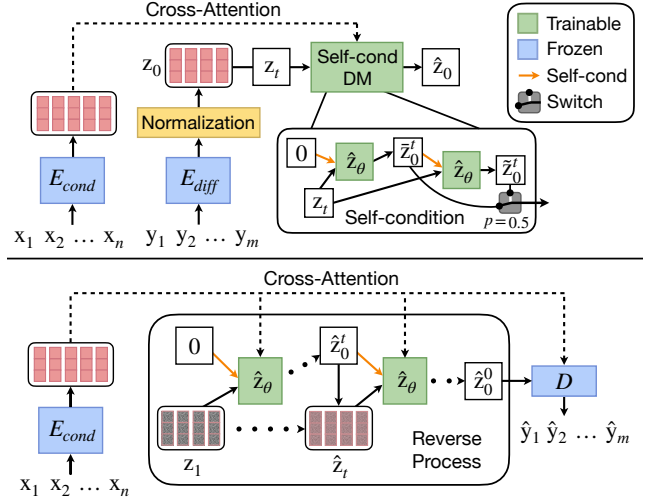


Figure 1: Overview of our framework design for conditional generation. Top is the training process, bottom is the generation process.

of short texts. After encoding the text, the encodings of all special tokens are replaced by their corresponding embeddings and padding encodings are replaced with zeros. This is necessary because diffusion model does not use an attention mask during training, which means that the reconstruction loss is calculated for both text and special tokens. However, special token encodings usually contain meaningless for diffusion model values. Therefore, minimization of reconstruction loss for these encodings only harms the training process. Embeddings of special tokens, on the other hand, only contain information about the token itself and the diffusion model recovers them much easier. During training we do not update the weights of the encoder in order to keep the approach simple.

### 5.2 Decoder, $D$

The decoder  $D$  is required to convert latent variables generated by diffusion model into textual output. Although a basic linear decoder can effectively reconstruct tokens with high accuracy, we employ the BERT-type (Devlin et al. 2019) architecture for the decoder to provide it with the ability to capture context information and rectify potential mistakes originating from the diffusion model. Note that we do not use an AR decoder on purpose so as not to transfer the limitations of AR language models to the diffusion model.

We train the decoder independently of the diffusion model using the following objective

$$-\mathbb{E} \log p_D(y | Cor(z_0)) \rightarrow \min, \quad (2)$$

where  $Cor(z_0)$  is a corrupted latent variable extracted from the diffusion encoder. Corruption is needed to expand the decoder training data domain and make it robust to distribution mismatch between text encodings  $z_0$  and latents  $\hat{z}_0$  generated by the diffusion model. This mismatch might arise

due to the accumulation of errors during the denoising process. Its presence is especially evident for special tokens, which always have the same fixed representations in  $z_0$ . By default, we take  $Cor(z_0)$  to be  $z_t$  with randomly sampled  $t \in [0, 0.15]$ . We use the diffusion’s noise scheduler to calculate  $z_t$ .

### 5.3 Diffusion model, $\hat{z}_\theta$

The diffusion model consists of 12 BERT layers and it is trained to predict the original latent  $z_0$  given its noisy version  $z_t$  and a timestep  $t$  by minimizing the objective (1). We provide the model with timestep value by adding timestep embedding to the hidden state vectors on each layer.

We train the diffusion model using the variance preserving scheme, discussed in (Song et al. 2021). To achieve zero mean and unit variance we normalize the latent variables  $z_0$  coordinate-wise, using the statistics from the training set.

**Noise scheduler** We adopt the noise scheduler from (Hoogeboom, Heek, and Salimans 2023) and use the following equation for  $\alpha_t$ :

$$\alpha_t = \frac{1}{1 + \tan(t\pi/2)^2 \cdot d^2}, \quad (3)$$

where  $d$  is a hyperparameter controlling the rate at which noise is introduced into the system. We set  $d = 9$  by default, which corresponds to a significantly higher noise addition rate than what is used in all common noise schedulers. We further refer to our scheduler as *tan-d* noise scheduler.

**Self-condition** Following the previous approaches (Lovelace et al. 2023; Strudel et al. 2022) we incorporate self-conditioning into the diffusion model. In order to make the model utilize the data sample estimation from the previous generation step, we modify the training procedure.

According to (Chen, Zhang, and Hinton 2023) we design the training process to emulate the inference behavior. On each training iteration with the probability  $p = 0.5$  the prediction is computed with the self-conditioning set to zero  $\tilde{z}_0^t = z_\theta(z_t, t, 0)$ . And, with probability  $(1 - p) = 0.5$  we first calculate  $\tilde{z}_0^t = z_\theta(z_t, t, 0)$  and then use it as an estimation of the data sample to obtain a second prediction  $\tilde{z}_0^t = z_\theta(z_t, t, \text{SG}(\tilde{z}_0^t))$ , where SG is the stop-gradient function that does not allow the gradient to flow through  $\tilde{z}_0^t$ . The diffusion model is optimized using the output  $\tilde{z}_0^t$  in the former scenario and  $\tilde{z}_0^t$  in the latter. This training strategy allows the model to accurately approximate  $z_0$  both with and without self-conditioning. We implement self-conditioning in a same manner as conditioning on timestep. For each diffusion model layer we pass the data estimation through a single linear layer and add it to the hidden state vectors.

### 5.4 Generation process

The generation process is illustrated on the Figure 1 (bottom). To generate text in the inference phase, we start with a random Gaussian sample and denoise it in  $T$  steps using the Euler solver. At each step, we apply self-conditioning and, because of it, use a small number of steps – 50 by default.

For the conditional generation we keep the framework design similar to unconditional generation. The only difference

is that we add condition encoder to process the input text and provide both diffusion model and decoder with its output via cross-attention. We also add classifier-free guidance with coefficient 0.5, because it slightly improves quality. Implementation details can be found in Appendix E.

## 6 Datasets

To evaluate the performance of our diffusion models we use five datasets in English language. Two of them are unconditional: **ROCStories** and **Wikipedia**, and three are conditional: **QQP**, **XSum** and **Wiki-Auto**. The **ROCStories** (Mostafazadeh et al. 2016) dataset contains 98k five-sentence commonsense fictional stories, that capture causal and temporal relations between daily events. The **Wikipedia** dataset is obtained from the ROOTS corpus (Laurençon et al. 2023), it is a collection of over 2 million cleaned articles from the Wikipedia platform. The subset of **QQP** (Chen et al. 2017) dataset, proposed in (Gong et al. 2023), consists of 144k question pairs from the Quora platform that are paraphrases of each other. The **XSum** (Narayan, Cohen, and Lapata 2018) dataset is used for summarization problem and it contains 204k BBC articles, which are provided as document and summary pairs. The **Wiki-Auto** (Jiang et al. 2020) dataset consists of aligned sentences from complex and simplified Wikipedia<sup>1</sup>. The detailed statistics for each dataset can be found in Appendix F.

## 7 Empirical Analysis

In this section, we evaluate the components of our framework on the **ROCStories** and **Wikipedia** datasets. To simplify the setup, we only consider unconditional generation. In Section 8, we demonstrate that our findings can be successfully transferred to the conditional generation problems. In this section, we do not compare our method with others. The comparison with the GPT-2 (Radford et al. 2019) on unconditional generation is presented in Appendix J.

### 7.1 Evaluation Metrics

We follow the model evaluation scheme from the (Lovelace et al. 2023). To evaluate the quality of our model we use **Perplexity (ppl)**, calculated with GPT-2 Large (Radford et al. 2019). To measure the diversity of the generated text we utilize the diversity metric proposed in (Su et al. 2022). We calculate it as  $\text{div}(y) = \prod_{n=2}^4 \frac{\# \text{ of unique n-grams in } y}{\# \text{ of n-grams in } y}$ , where  $y$  is a set of generated texts. To ensure that the model does not reproduce the training dataset during the generation we evaluate the **Memorization (mem)**. We calculate it as the proportion of generated 4-grams that are found in the training set. As Perplexity tends to be small for the texts with repetitions, we also measure **MAUVE Score** (Pillutla et al. 2021) to estimate the quality of text. MAUVE is a language model-based metric that measures the distance between the distributions of generated and reference texts using divergence frontiers. We leave all MAUVE hyperparameters at the default values presented in the original paper.

<sup>1</sup>All the datasets we use in this work are publicly available under a creative commons or an open source license.

Encoder	ppl ↓	mem ↓	div ↑	mauve ↑
<b>ROCStories</b>				
BERT emb	48.9 <sub>.36</sub>	.371 <sub>.003</sub>	.324 <sub>.002</sub>	.600 <sub>.016</sub>
BERT	29.1 <sub>.89</sub>	.453 <sub>.003</sub>	.295 <sub>.002</sub>	<b>.762</b> <sub>.043</sub>
RoBERTa	<b>28.3</b> <sub>.33</sub>	.443 <sub>.003</sub>	.302 <sub>.002</sub>	.647 <sub>.019</sub>
T5	31.3 <sub>.54</sub>	.427 <sub>.003</sub>	.312 <sub>.004</sub>	.706 <sub>.024</sub>
BART	34.1 <sub>.52</sub>	.441 <sub>.006</sub>	.299 <sub>.005</sub>	.705 <sub>.030</sub>
Source text	21.7	.365	.403	.876
<b>Wikipedia</b>				
BERT emb	156.1 <sub>1.8</sub>	.263 <sub>.004</sub>	.517 <sub>.002</sub>	.378 <sub>.055</sub>
BERT	<b>104.4</b> <sub>2.1</sub>	.286 <sub>.002</sub>	.504 <sub>.003</sub>	<b>.874</b> <sub>.011</sub>
Source text	37.3	.122	.615	.957

Table 1: Quality of diffusion models trained with different diffusion encoders.

To calculate all the metrics, we generate 1000 texts. For MAUVE, we sample 1000 reference texts from the test set. We repeat this procedure 5 times and report the mean and standard deviation of the results in the  $\text{mean}_{\text{std}}$  notation.

## 7.2 Model setup

Unless otherwise stated, the training of **TEncDM** is performed within the latent space of BERT encodings. A three-layer transformer is employed for the decoder, which is trained to reconstruct  $z_0$  from  $z_t$ , where  $t \in U[0, 0.15]$ . A comprehensive analysis of various decoder modifications is presented in this section and Appendix B. The diffusion model is a 12-layer transformer with a dimensionality of 768. We train it with *tan-9* noise scheduler.

**Effect of Diffusion Encoder** We compare latent spaces of BERT (Devlin et al. 2019), RoBERTa (Liu et al. 2019), BART (Lewis et al. 2020) and T5 (Raffel et al. 2020) encoders, as well as BERT embeddings, to ascertain the optimal choice for the diffusion model. All encoders have approximately the same size of 100M parameters. In this experiment, we train diffusion models with the same set of hyperparameters across all diffusion encoders. We train the decoders according to the scheme described in Section 7.2. The results of this comparison are presented in Table 1 and they show a clear advantage of the latent space derived from BERT encodings on ROCStories dataset. Furthermore, the quality of all encoders is superior to that of BERT embeddings. A better **div** and **mem** for embeddings can be explained by the presence of words in the corpus that do not align with the context. The text samples are presented in Appendix K. This confirms our hypothesis that encodings are better suited for the training of a diffusion model. We discuss the drop in **mauve** for RoBERTa in Appendix G.

**Effect of Decoder** To confirm the hypothesis about the importance of the decoder architecture and its training scheme, we compare an MLP decoder consisting of two linear layers with a 3-layer transformer. The latter is able to extract contextual information for each token. We corrupt the decoder input  $z_0$  by transforming it into  $z_t$ , using the diffusion forward process with  $t \in U[0, 0.15]$ . We choose this method, because it brings the decoder input closer to the diffusion

Decoder	ppl ↓	mem ↓	div ↑	mauve ↑
<b>ROCStories</b>				
MLP	39.7 <sub>3.38</sub>	.444 <sub>.002</sub>	.297 <sub>.004</sub>	.716 <sub>.074</sub>
+ <i>Cor</i> ( $z_0$ )	31.2 <sub>.33</sub>	.448 <sub>.002</sub>	.293 <sub>.003</sub>	.739 <sub>.051</sub>
Transformer	34.2 <sub>.29</sub>	.445 <sub>.001</sub>	.295 <sub>.003</sub>	.714 <sub>.037</sub>
+ <i>Cor</i> ( $z_0$ )	<b>29.1</b> <sub>.89</sub>	.453 <sub>.003</sub>	.295 <sub>.002</sub>	<b>.762</b> <sub>.043</sub>
<b>Wikipedia</b>				
Transformer	180.6 <sub>3.2</sub>	.261 <sub>.001</sub>	.511 <sub>.001</sub>	.526 <sub>.025</sub>
+ <i>Cor</i> ( $z_0$ )	<b>104.4</b> <sub>2.1</sub>	.286 <sub>.002</sub>	.504 <sub>.003</sub>	<b>.874</b> <sub>.011</sub>

Table 2: Comparison of decoders for encoding-based diffusion model.

output. A more detailed analysis of corruption techniques is presented in the Appendix B. To keep the experiment fair, we apply all decoders to the same generated latents. The results of the experiment are shown in Table 2. The MLP decoder without corruption achieves the lowest text quality in terms of **perplexity**, but comparable by **mauve** with Transformer without corruption. However, it is challenging to make meaningful comparisons between decoders by **mauve** due to the significant variance. From this experiment, we can conclude that corruption of the latent helps to improve the quality for both datasets. At the same time, the incorporation of contextual information into the decoder lead to the best result.

**Effect of self-conditioning** We conduct a series of experiments to understand how self-conditioning (SC) affects the denoising process. In Figure 2, we compare the quality of the models with and without SC for different number of denoising steps on the ROCStories dataset. The results show that while the quality of the model without SC increases as the number of steps increases, the quality of the model with SC reaches a maximum at a value of 50 steps in terms of **mauve**, and then it starts to drop. Nevertheless, at the highest point the model with SC surpasses the model without it according to both **mauve** and **perplexity**.

We explain this drop in generation quality with mismatch between diffusion model inputs at train and inference stages. To confirm our hypothesis, we calculated the mean-squared norm (*magnitude*) of the values of each latent  $\hat{z}_0^t$  in a mini-batch predicted by the diffusion model during generation (i.e.  $\frac{1}{N \cdot d \cdot m} \|\hat{z}_0^t\|_2^2$ , where  $N$  is a batch size,  $d$  is a dimension and  $m$  is a sequence length). We plot this magnitude with respect to timestep for generations with different number of steps as well as for the predictions  $\bar{z}_0^t$  from the training stage. The results for the ROCStories dataset are presented in Figure 3. They indicate that self-conditioning significantly increases the prediction magnitude as the number of steps increases. This can be explained by the following: during training, the model learns to use self-conditioning to approximate  $z_0$  more accurately. Consequently, self-conditioning increases the model’s confidence, which is directly related to prediction magnitude. During the generation process, the model takes its own prediction, which has an increased magnitude, as an input at each step and increases it further. Therefore, the increase in magnitude depends directly on




Figure 2: Generation quality of diffusion models with and without self-conditioning on ROCStories dataset.




Figure 3: Prediction magnitudes for generation processes with different amount of steps on ROCStories dataset.

the number of generation steps. Eventually, this leads to a mismatch between the predictions fed into the model during training and generation.

In the Appendix C, we provide a more detailed discussion of this phenomenon and show that the same behavior is observed in the larger Wikipedia dataset. It is worth noting that the smallest mismatch is observed for the trajectory of 50 generation steps, which corresponds to the best quality.

**Effect of Noise scheduler** We compare our noise scheduler  $\tan-d$  with previously used  $\cosine$  and  $\sqrt{t}$  (visualized in Appendix D) and present the quantitative results in Table 3. We use the same decoder and optimal amount of generation steps for each scheduler. In Figure 4, we evaluate the difficulty of recovering a data sample from noised latent  $z_t$  for diffusion model trained with different noise schedulers. We measure the reconstruction loss  $\frac{1}{N \cdot d \cdot m} \|z_0 - \hat{z}_0^t\|_2^2$  and accuracy of token prediction for every timestep.

While the  $\sqrt{t}$  noise scheduler adds significantly larger amount of noise in the initial timesteps than  $\cosine$  one, the rate of noise addition decreases for the subsequent timesteps. For both schedulers, the denoising task becomes insufficiently hard for the most timesteps, which leads to a de-




Figure 4: Reconstruction loss and reconstruction accuracy of diffusion models trained with different noise schedulers on ROCStories dataset.

Noise Scheduler	ppl ↓	mem ↓	div ↑	mauve ↑
<b>ROCStories</b>				
cosine	393.2 <sub>127.6</sub>	.262 <sub>.004</sub>	.474 <sub>.006</sub>	.098 <sub>.011</sub>
sqrt	127.2 <sub>29.3</sub>	.264 <sub>.004</sub>	.434 <sub>.004</sub>	.364 <sub>.041</sub>
tan-7	<b>27.1</b> <sub>.31</sub>	.455 <sub>.004</sub>	.286 <sub>.001</sub>	.730 <sub>.026</sub>
tan-9	29.1 <sub>.89</sub>	<b>.453</b> <sub>.003</sub>	<b>.295</b> <sub>.002</sub>	<b>.762</b> <sub>.043</sub>
tan-11	33.4 <sub>2.4</sub>	.464 <sub>.006</sub>	.279 <sub>.004</sub>	.730 <sub>.038</sub>
<b>Wikipedia</b>				
sqrt	364.0 <sub>6.5</sub>	.139 <sub>.001</sub>	.664 <sub>.004</sub>	.325 <sub>.037</sub>
tan-9	<b>104.4</b> <sub>2.1</sub>	.286 <sub>.002</sub>	.504 <sub>.003</sub>	<b>.874</b> <sub>.011</sub>

Table 3: Quality of diffusion models trained with different noise schedulers.

crease in their contribution to the generation process. This can be seen from the reconstruction accuracy. In contrast,  $\tan-d$  noise scheduler adds more noise consistently across all timesteps, leading to a more challenging training task and improved generation performance.

Based on these observations, we conclude that in order to improve the efficiency of the denoising process, it is essential to increase the amount of added noise within all timesteps. However, it is important to strike a balance as adding excessive noise can negatively impact performance. In our experiments,  $\tan-9$  yielded results that were marginally superior to all other schedulers in terms of **mauve**, with slightly better **memorization** and **diversity**, while lagging behind  $\tan-7$  in terms of **perplexity**.

As a rule of thumb, the noise schedule should be such that the diffusion model recovers approximately the same amount of information at each timestep. Otherwise, some of them will not contribute to the denoising process enough.

## 8 Seq2Seq Experiments

We are conducting experiments to evaluate the effectiveness of the proposed components for text diffusion generation on language model encodings on three different tasks: paraphrasing (QQP), summarization (XSum), and text simplification (Wiki-Auto).

Method	QQP			XSum			Wiki-Auto		
	R-L ↑	BS ↑	B-4 ↑	R-1/2/L ↑	BS ↑	R-L ↑	BS ↑	B-4 ↑	
DiffuSeq*	52.7	82.4	—	18.9 / 1.3 / 13.6	46.8	—	79.1	26.1	
SeqDiffuSeq*†	—	82.9	23.3	14.1 / 1.1 / 11.4	58.4	—	<b>82.1</b>	37.1	
GENIE*	—	—	—	29.3 / 8.3 / 21.9	—	—	—	—	
AR-Diffusion†	54.9	81.4	<b>31.2</b>	27.1 / 6.4 / 20.8	59.7	54.9	81.3	32.7	
TEncDM (BERT)	56.4	82.4	30.2	31.5 / 10.0 / 24.9	68.2	<b>58.1</b>	80.5	<b>41.6</b>	
TEncDM (T5)	<b>57.3</b>	<b>83.8</b>	30.7	33.4 / 11.4 / 26.8	<b>70.1</b>	57.7	81.2	<b>41.6</b>	
TEncDM (RoBERTa)	55.8	82.4	30.0	<b>33.7 / 11.9 / 27.1</b>	69.8	57.9	81.0	40.5	
GPT2-small FT*	52.1	82.5	19.8	—	—	54.6	80.2	30.8	
Transformer-base*	57.5	83.8	27.2	30.5 / 10.4 / 24.2	—	49.1	73.8	26.9	
FLAN-T5-base*	52.3	83.2	—	34.6 / 12.9 / 27.2	72.7	—	—	—	

Table 4: Seq2Seq evaluation results of Diffusion and AR methods on QQP, XSum and Wiki-Auto datasets. We calculate **ROUGE-1/2/L** (**R-1/2/L**), **BERTScore** (**BS**) and **BLEU-4** (**B-4**). All results taken from other papers are marked with \*. DiffuSeq and SeqDiffuSeq results were taken from their respective publications. Results for AR models were taken from (Gong et al. 2023; Wu et al. 2023; Lovelace et al. 2023). Additionally, we trained AR-Diffusion and SeqDiffuSeq on previously unreported datasets, using the code from the corresponding papers (marked as †).

**Baselines** We include two groups of baselines in comparison. The first group comprises of popular diffusion methods: DiffuSeq (Gong et al. 2023), SeqDiffuSeq (Yuan et al. 2024), GENIE (Lin et al. 2023), AR-diffusion (Wu et al. 2023). We focus only on non-autoregressive diffusion models trained in the latent space of embeddings or encodings. Besides, we compare TEncDM to classical AR baselines: Transformer (Vaswani et al. 2017a), FLAN-T5-base (Chung et al. 2024) and finetuned GPT-2-small (Radford et al. 2019).

**Metrics** For evaluation of paraphrasing and simplification tasks, we adopt the setting of SeqDiffuSeq (Yuan et al. 2024) and calculate ROUGE-L (Lin 2004), BERTScore (Zhang et al. 2020) and BLEU-4. In addition, we follow the approach of (Wu et al. 2023) and report ROUGE-1/2/L for summarization task.

**Results** Table 4 presents a comprehensive comparison of our approach against existing methods across three datasets. Results for DiffuSeq, and SeqDiffuSeq were sourced from their respective papers (Wu et al. 2023; Yuan et al. 2024). Results for GENIE were taken from the (Wu et al. 2023), as the original paper used ground truth labels to select the best generated text, which introduces unfairness. Additionally, we trained both AR-Diffusion and SeqDiffuSeq on previously unreported datasets, using the code from the corresponding papers. Results for AR models were taken from (Gong et al. 2023; Wu et al. 2023; Lovelace et al. 2023).

We experiment with three encoders: BERT-base, T5-base, RoBERTa-base to investigate their efficacy on conditional tasks. We use the same encoder for  $E_{diff}$  and  $E_{cond}$ . Our findings indicate that all three encoders demonstrate effectiveness across the tasks, achieving comparable performance levels. However, no single encoder outperforms the others across all tasks. T5-base excels in question paraphrasing (QQP), while RoBERTa-base demonstrates superior performance on summarization (XSum) and on text simplification (Wiki-Auto) all encoders exhibit similar performance.

A comparison with other methods clearly shows that **TEncDM**, using any of the tested encoder models, outper-

forms popular diffusion embedding-based approaches. Furthermore, **TEncDM** with an optimal encoder achieves comparable performance to three AR baselines on QQP and XSum and outperforms them on Wiki-Auto.

## 9 Conclusion

In this work, we explore key details of the diffusion pipeline for text generation. We propose **TEncDM** which trains the diffusion model in the latent space of the language model encoding. Unlike embeddings, they contain contextual information which helps diffusion model to recover latents. To improve text generation performance, we analyse the effect of self-conditioning and conclude that it increases the magnitude of the model’s predictions, which allows to reduce the number of generation steps. In addition, we propose an efficient decoder and noise scheduler that boost the generation quality. Extensive ablation on ROCStories and Wikipedia datasets demonstrates the impact of the proposed design choices. Finally, **TEncDM** outperforms recent diffusion models and some classical autoregressive methods in downstream task experiments.

## 10 Limitations

There are three limitations that warrant further investigation. First, the quality of the model can be improved by training diffusion encoder, decoder and denoising model simultaneously. However, we avoid doing so in order to avoid overcomplicating the approach. Second, the samples from the latent space have a high dimensionality that depends on the sequence length, so the training of our method slows down significantly as the length increases. This problem can probably be overcome by training the autoencoder to map a text into a fixed-dimensional latent space, which is a great direction for further research. Third, as different diffusion encoders works better for different tasks, it is necessary to find the best one for each task.

## Acknowledgments

The paper was prepared within the framework of the HSE University Basic Research Program and was supported in part through computational resources of HPC facilities at HSE University.

## References

- Austin, J.; Johnson, D. D.; Ho, J.; Tarlow, D.; and van den Berg, R. 2021. Structured Denoising Diffusion Models in Discrete State-Spaces. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 17981–17993. Curran Associates, Inc.
- Blattmann, A.; Dockhorn, T.; Kulal, S.; Mendelevitch, D.; Kilian, M.; Lorenz, D.; Levi, Y.; English, Z.; Voleti, V.; Letts, A.; Jampani, V.; and Rombach, R. 2023. Stable Video Diffusion: Scaling Latent Video Diffusion Models to Large Datasets. arXiv:2311.15127.
- Chen, T.; Zhang, R.; and Hinton, G. 2023. Analog Bits: Generating Discrete Data using Diffusion Models with Self-Conditioning. arXiv:2208.04202.
- Chen, Z.; Zhang, H.; Zhang, X.; and Zhao, L. 2017. Quora Question Pairs.
- Chung, H. W.; Hou, L.; Longpre, S.; Zoph, B.; Tay, Y.; Fedus, W.; Li, Y.; Wang, X.; Dehghani, M.; Brahma, S.; et al. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70): 1–53.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Burstein, J.; Doran, C.; and Solorio, T., eds., *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; and et al. 2024. The Llama 3 Herd of Models. arXiv:2407.21783.
- Evans, Z.; Carr, C.; Taylor, J.; Hawley, S. H.; and Pons, J. 2024. Fast Timing-Conditioned Latent Audio Diffusion. arXiv:2402.04825.
- Gao, Z.; Guo, J.; Tan, X.; Zhu, Y.; Zhang, F.; Bian, J.; and Xu, L. 2024. Empowering Diffusion Models on the Embedding Space for Text Generation. In Duh, K.; Gomez, H.; and Bethard, S., eds., *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 4664–4683. Mexico City, Mexico: Association for Computational Linguistics.
- Gong, S.; Li, M.; Feng, J.; Wu, Z.; and Kong, L. 2023. DiffuSeq: Sequence to Sequence Text Generation with Diffusion Models. In *The Eleventh International Conference on Learning Representations*.
- Han, X.; Kumar, S.; and Tsvetkov, Y. 2023. SSD-LM: Semi-autoregressive Simplex-based Diffusion Language Model for Text Generation and Modular Control. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 11575–11596. Toronto, Canada: Association for Computational Linguistics.
- Hoogeboom, E.; Heek, J.; and Salimans, T. 2023. Simple diffusion: end-to-end diffusion for high resolution images. In *Proceedings of the 40th International Conference on Machine Learning, ICML’23*. OpenReview.net.
- Hoogeboom, E.; Nielsen, D.; Jaini, P.; Forré, P.; and Welling, M. 2021. Argmax Flows and Multinomial Diffusion: Learning Categorical Distributions. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 12454–12465. Curran Associates, Inc.
- Jiang, C.; Maddela, M.; Lan, W.; Zhong, Y.; and Xu, W. 2020. Neural CRF Model for Sentence Alignment in Text Simplification. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- Karimi Mahabadi, R.; Ivison, H.; Tae, J.; Henderson, J.; Beltagy, I.; Peters, M.; and Cohan, A. 2024. TESS: Text-to-Text Self-Conditioned Simplex Diffusion. In Graham, Y.; and Purver, M., eds., *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2347–2361. St. Julian’s, Malta: Association for Computational Linguistics.
- Laurençon, H.; Saulnier, L.; Wang, T.; Akiki, C.; del Moral, A. V.; Scao, T. L.; Werra, L. V.; Mou, C.; Ponferrada, E. G.; Nguyen, H.; Frohberg, J.; Šaško, M.; Lhoest, Q.; McMillan-Major, A.; Dupont, G.; Biderman, S.; Rogers, A.; allal, L. B.; Toni, F. D.; Pistilli, G.; Nguyen, O.; Nikpoor, S.; Masoud, M.; Colombo, P.; de la Rosa, J.; Villegas, P.; Thrush, T.; Longpre, S.; Nagel, S.; Weber, L.; Muñoz, M.; Zhu, J.; Strien, D. V.; Alyafeai, Z.; Almubarak, K.; Vu, M. C.; Gonzalez-Dios, I.; Soroa, A.; Lo, K.; Dey, M.; Suarez, P. O.; Gokaslan, A.; Bose, S.; Adelani, D.; Phan, L.; Tran, H.; Yu, I.; Pai, S.; Chim, J.; Lepercq, V.; Ilic, S.; Mitchell, M.; Luccioni, S. A.; and Jernite, Y. 2023. The BigScience ROOTS Corpus: A 1.6TB Composite Multilingual Dataset. arXiv:2303.03915.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 7871–7880. Online: Association for Computational Linguistics.
- Li, X.; Thickstun, J.; Gulrajani, I.; Liang, P. S.; and Hashimoto, T. B. 2022. Diffusion-LM Improves Controllable Text Generation. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 4328–4343. Curran Associates, Inc.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.

- Lin, Z.; Gong, Y.; Shen, Y.; Wu, T.; Fan, Z.; Lin, C.; Duan, N.; and Chen, W. 2023. Text generation with diffusion language models: a pre-training approach with continuous paragraph denoise. In *Proceedings of the 40th International Conference on Machine Learning, ICML'23*. JMLR.org.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. arXiv:1907.11692.
- Lovelace, J.; Kishore, V.; Wan, C.; Shekhtman, E.; and Weinberger, K. Q. 2023. Latent Diffusion for Language Generation. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems*, volume 36, 56998–57025. Curran Associates, Inc.
- Meng, C.; Rombach, R.; Gao, R.; Kingma, D.; Ermon, S.; Ho, J.; and Salimans, T. 2023. On Distillation of Guided Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14297–14306.
- Mostafazadeh, N.; Chambers, N.; He, X.; Parikh, D.; Batra, D.; Vanderwende, L.; Kohli, P.; and Allen, J. 2016. A Corpus and Cloze Evaluation for Deeper Understanding of Commonsense Stories. In Knight, K.; Nenkova, A.; and Rambow, O., eds., *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 839–849. San Diego, California: Association for Computational Linguistics.
- Narayan, S.; Cohen, S. B.; and Lapata, M. 2018. Don't Give Me the Details, Just the Summary! Topic-Aware Convolutional Neural Networks for Extreme Summarization. In Riloff, E.; Chiang, D.; Hockenmaier, J.; and Tsujii, J., eds., *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 1797–1807. Brussels, Belgium: Association for Computational Linguistics.
- OpenAI. 2023. GPT-4 Technical Report. *ArXiv*, abs/2303.08774.
- Pillutla, K.; Swayamdipta, S.; Zellers, R.; Thickstun, J.; Welleck, S.; Choi, Y.; and Harchaoui, Z. 2021. MAUVE: Measuring the Gap Between Neural Text and Human Text using Divergence Frontiers. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 4816–4828. Curran Associates, Inc.
- Podell, D.; English, Z.; Lacey, K.; Blattmann, A.; Dockhorn, T.; Müller, J.; Penna, J.; and Rombach, R. 2023. SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis. arXiv:2307.01952.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language Models are Unsupervised Multitask Learners. [https://d4mucfpsywv.cloudfront.net/better-language-models/language\\_models\\_are\\_unsupervised\\_multitask\\_learners.pdf](https://d4mucfpsywv.cloudfront.net/better-language-models/language_models_are_unsupervised_multitask_learners.pdf). OpenAI blog, p. 9.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research*, 21(140): 1–67.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10684–10695.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2021. Score-Based Generative Modeling through Stochastic Differential Equations. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Strudel, R.; Tallec, C.; Alché, F.; Du, Y.; Ganin, Y.; Mensch, A.; Grathwohl, W.; Savinov, N.; Dieleman, S.; Sifre, L.; and Leblond, R. 2022. Self-conditioned Embedding Diffusion for Text Generation. arXiv:2211.04236.
- Su, Y.; Lan, T.; Wang, Y.; Yogatama, D.; Kong, L.; and Collier, N. 2022. A Contrastive Framework for Neural Text Generation. In Oh, A. H.; Agarwal, A.; Belgrave, D.; and Cho, K., eds., *Advances in Neural Information Processing Systems*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017a. Attention is all you need. *Advances in neural information processing systems*, 30.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017b. Attention is All you Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Wu, T.; Fan, Z.; Liu, X.; Zheng, H.-T.; Gong, Y.; yelong shen; Jiao, J.; Li, J.; zhongyu wei; Guo, J.; Duan, N.; and Chen, W. 2023. AR-Diffusion: Auto-Regressive Diffusion Model for Text Generation. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Ye, J.; Zheng, Z.; Bao, Y.; Qian, L.; and Wang, M. 2024. DI-NOISER: Diffused Conditional Sequence Learning by Manipulating Noises. arXiv:2302.10025.
- Yuan, H.; Yuan, Z.; Tan, C.; Huang, F.; and Huang, S. 2024. Text Diffusion Model with Encoder-Decoder Transformers for Sequence-to-Sequence Generation. In Duh, K.; Gomez, H.; and Bethard, S., eds., *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 22–39. Mexico City, Mexico: Association for Computational Linguistics.
- Zhang, T.; Kishore, V.; Wu, F.; Weinberger, K. Q.; and Artzi, Y. 2020. BERTScore: Evaluating Text Generation with BERT. arXiv:1904.09675.
- Zhang, Y.; Gu, J.; Wu, Z.; Zhai, S.; Susskind, J. M.; and Jaitly, N. 2023. PLANNER: Generating Diversified Paragraph via Latent Language Diffusion Model. In *Thirty-seventh Conference on Neural Information Processing Systems*.