

tendency to generate consecutive repetitions at the word, phrase and sentence levels. For example, with repetitive prompt, the model can enter an *absorbing state* where it produces repetitive outputs with higher and higher confidence by *self-reinforcing* the pattern (Xu et al., 2022) (Fig. 1). Through our experiments, we have observed that such degeneration is more prevalent in *open-ended* tasks that allow the model greater freedom for creativity. Even for large language models, the generation can drift away from the desired semantics, especially when the model is poorly prompted or has high initial probabilities (Xu et al., 2022).

Why does using maximum likelihood decoding lead to repetitions during generation which is significantly different from the training data distribution? One possible explanation for this is “exposure bias” (Bengio et al., 2015) arising from the discrepancy between the training and inference phases in the teacher forcing training strategy. During training phase, the model focuses only on predicting the next token. However, during inference, predicting the next token alone can be myopic because the model may not have enough foresight to anticipate its impact on future generation. This can also be seen as the “distribution shift” issue of behavior cloning (de Haan et al., 2019), where the model is trained to mimic the expert’s actions on the states encountered by the expert in the training data. However, small differences between the model and the expert can compound over multiple steps, leading the model to states it never encountered during training, rendering unreliable and undesirable predictions.

Although many approaches have been proposed to address this issue, such as adversarial models (Yu et al., 2017; Lamb et al., 2016; Zhang et al., 2017), reinforcement learning (Li et al., 2016b) or repetition penalties (Xu et al., 2022), they attempt to improve the *global* aspects of the generation by making *local* adjustments which still follow the autoregressive generation recipe. Diffusion models provide an alternative solution – the model can revisit and revise its output iteratively, potentially rendering more global control of the generation in a non-autoregressive manner. However, these text diffusion models can generate less fluent text compared to autoregressive ones (Gong et al., 2023). Also, when generating long text, the diffusion process involves multiple passes of the underlying denoising model over a long generation length, making it computationally expensive. The discrete nature of text also presents a challenge for diffusion models, which can suffer from “rounding errors” when converting between the text token and its embedding (Li et al., 2022; Lin et al., 2022).

Instead of performing diffusion on the original text or the corresponding word embeddings, we propose to apply diffusion techniques to the latent semantic space (Rombach et al., 2022; Lovelace et al., 2022). To achieve this, we learn a fixed number of continuous semantic tokens that encode salient information at the paragraph level. These tokens can then be used to reconstruct the original text. The latent diffusion can be additionally conditioned on an external signal to generate the semantic tokens. Finally, a decoder maps the obtained semantic tokens back to the raw text space. This process combines a non-autoregressive semantic diffusion approach with an autoregressive decoding technique. The semantic diffusion process handles the “planning”, enabling the modification of semantics in a coarse-to-fine manner, while the decoder handles the “decoding” by translating the semantics into raw text, with less flexibility in controlling the meaning. We call our proposed method PLANNER (Paragraph-level Diffusion model for Embedding Representation).

Our contributions include: (i) We propose a latent semantic diffusion model for paragraphs that incorporates both non-autoregressive semantic diffusion and autoregressive generation. This allows us to generate fluent text while being able to exercise global control inherited from a diffusion model. (ii) We study the essential requirements for a good latent space for paragraph diffusion models. (iii) We evaluate the effectiveness of our proposed method on various conditional generation tasks. Thanks to the iterative refinement of denoising diffusion, our method enjoys less repetitive and more diverse generation, while maintaining good fluency and relevancy, comparing with autoregressive and text diffusion baselines (Li et al., 2022; Lin et al., 2022).

2 Preliminary

Diffusion Probabilistic Models The standard diffusion model (DM) (Ho et al., 2020; Song & Ermon, 2019) learns the data distribution $p(x)$ by gradually denoising a normally distributed variable in a Markov chain of length T . The diffusion process can be viewed as a continuous-time stochastic process (Song et al., 2021b; Kingma et al., 2021) where the initial data point $\mathbf{x} \in \mathbb{R}^N$ is progressively corrupted by noise according to a predefined signal-noise schedule $\{\alpha_t, \sigma_t\}$, resulting in time-

dependent corrupted data $\{\mathbf{x}_t | t \in [0, 1], \mathbf{x}_0 = \mathbf{x}\}$. The transition distribution is given by:

$$q(\mathbf{x}_t | \mathbf{x}_s) = \mathcal{N}(\mathbf{x}_t; \alpha_t |_s \mathbf{x}_s, \sigma_t^2 |_s I), \quad (1)$$

where $\alpha_t |_s = \alpha_t / \alpha_s$, $\sigma_t^2 |_s = \sigma_t^2 - \alpha_t^2 |_s \sigma_s^2$, and $s < t$. When $\mathbf{x}_s = \mathbf{x}$, the marginal distribution $q(\mathbf{x}_t | \mathbf{x})$ is given as $q(\mathbf{x}_t | \mathbf{x}) = \mathcal{N}(\mathbf{x}_t; \alpha_t \mathbf{x}, \sigma_t^2 I)$. The diffusion model relies on a parametric function θ optimized to reverse the diffusion process by denoising \mathbf{x}_t to the clean input \mathbf{x} . The model is trained using a weighted reconstruction loss:

$$\mathcal{L}(\theta) = \mathbb{E}_{\mathbf{x}_t \sim q(\mathbf{x}_t | \mathbf{x}), t \sim [0, 1]} [\omega_t \cdot \|\mathbf{F}_\theta(\mathbf{x}_t, t) - \mathbf{x}\|_2^2], \quad (2)$$

where $\omega_t = \alpha_t^2 / \sigma_t^2$, (*s.t.* $\alpha_t^2 + \sigma_t^2 = 1$) is the signal-to-noise-ratio (SNR) and $\mathbf{F}_\theta(\cdot)$ denotes the backbone denoising function. Sampling from the learned model can be performed using either ancestral sampling (DDPM) (Ho et al., 2020) or a deterministic DDIM sampler (Song et al., 2021a). While the DM is capable of generating high-quality samples, the fact that the corrupted data \mathbf{x}_t shares the same space as the input \mathbf{x} results in inefficient training (Jing et al., 2022) and difficulty in learning abstract and semantically meaningful latent spaces (Preechakul et al., 2022).

Latent Diffusion Models To improve the efficiency, the Latent Diffusion Model (LDM) (Rombach et al., 2022) introduces an explicit separation between the compressive and generative learning phases of training diffusion models. It employs an autoencoding model consisting of an encoder $\mathcal{E}(\cdot)$ and a decoder $\mathcal{D}(\cdot)$ to learn a low-dimensional latent space that is perceptually equivalent to the image space when decoded, but with reduced computational complexity, while retaining the perceptual quality of generated samples. The reweighted objective for training LDM is given by:

$$\mathcal{L}(\theta) = \mathbb{E}_{\mathbf{z}_t \sim q(\mathbf{z}_t | \mathbf{z}), \mathbf{z} = \mathcal{E}(\mathbf{x}), t \sim [0, 1]} [\omega_t \cdot \|\mathbf{F}_\theta(\mathbf{z}_t, t) - \mathbf{z}\|_2^2], \quad (3)$$

where \mathbf{z} is obtained from \mathcal{E} during training. The generated \mathbf{z} can be decoded to image using \mathcal{D} .

3 Related Work

Text diffusion models Early attempts on using diffusion models for discrete data used a noising processes which masked or randomly mutated the discrete tokens (Austin et al., 2021; Hoogeboom et al., 2021). Recently, Diff-LM (Li et al., 2022) and DiffuSeq (Gong et al., 2023) have instead used a continuous token embedding space, converting the continuous token embeddings to text via "rounding". Analog Bits (Chen et al., 2022) converts raw text into a set of bits and models them as analog bits with a continuous diffusion model. (Lovelace et al., 2022) performed diffusion model on the contextualized BART embeddings rather than on the word embedding space. (Zhu et al., 2022) has applied text diffusion to image-captioning and achieved good performance.

However, existing text diffusion models present several issues: (i) The varying length of the input text necessitates the prediction of additional length or superfluous paddings, and (ii) token generation in parallel may result in disfluent text and/or frequent repetitions especially when the generation is long. We instead employ the diffusion model to learn paragraph embeddings that contain fewer fixed-sized tokens, which allows for computational benefits and improved fluency.

Text Variational Autoencoders Text VAEs (Bowman et al., 2016; Kim et al., 2018; Li et al., 2020) have been particularly useful for learning a smooth and interpretable representation space, as well as for generating diverse text. However, one of the challenges is the KL vanishing problem (Bowman et al., 2016), which results in the decoder disregarding the latent code sampled from the prior distribution during the inference stage. Our approach can be perceived as to address this issue by leveraging a more flexible prior distribution to ensure the codes can strongly influence the output text distribution.

4 PLANNER: A Language Diffusion Model on Paragraph Embeddings

We use latent diffusion to improve the diversity and fluency of paragraphs generated from the model. Our model comprises two parts (Fig. 2) - a paragraph embedder via variational autoencoder (VAE) that learns a meaningful and smooth latent space that corresponds to the original text space, and a diffusion model that generates latent codes corresponding to the semantics of longer text.

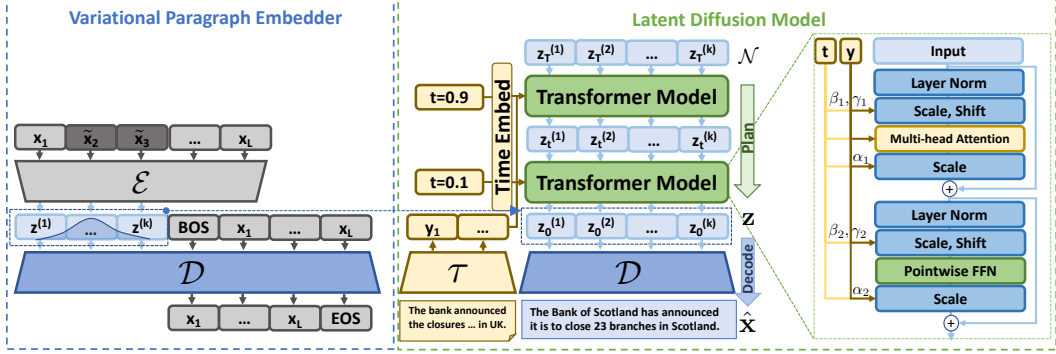


Figure 2: Model overview. Left: a variational paragraph embedder is learned to encode paragraph into a fixed amount of latent codes. Right: the latent diffusion model based on transformer block is applied to generate the latent codes. The decoder finally translates them into the text. (*BOS*: Begin of Sentence token, *EOS*: End of Sentence token)

4.1 Learning a Variational Paragraph Embedder

Instead of applying diffusion to tokens directly to generate long text, we propose to learn a set of latent codes $z = \{z^{(1)}, \dots, z^{(k)}\} \in \mathbb{R}^{k \times h}$, which we call *paragraph embeddings*, that capture the semantics in the target text (of length up to 512 tokens), where h denotes the embedding dimension. These paragraph embeddings have shorter length, such as $k = 16$, than the original text.

To obtain such embeddings z , we train a transformer-based encoder-decoder model. The architecture used for the autoencoder is shown in Fig. 2. The encoder \mathcal{E} and decoder \mathcal{D} construct a bidirectional mapping between the discrete data space and the latent code space. The paragraph embeddings z are extracted by taking the first k hidden state vectors of dimension h from the final layer of \mathcal{E} , which are fed into the initial steps of the decoder which is trained to reconstruct the original text. It’s worth noting that the paragraph embeddings share the same hidden dimension h as the word embeddings, and forming a manifold in the word embedding space. Pretrained BERT and GPT-2 models are used to initialize \mathcal{E} and \mathcal{D} , respectively. The manifold of the learned embeddings ideally possesses several desirable properties, including low *conversion error*, *local smoothness* and *distributional smoothness*.

Conversion error Ideally, the original input x can be perfectly reconstructed via $\hat{x} = \mathcal{D}(z)$, $z = \mathcal{E}(x)$, and modeling the lower-dimensional continuous space $p(z)$ is equivalent to modeling $p(x)$. However, in practice a loss of information can occur when converting raw text into paragraph embeddings or when doing the reverse. We assess the conversion loss by computing the BLEU score ($\text{BLEU}_{\text{clean}}$) between the input x and the reconstruction \hat{x} .

Local smoothness To generate target text that is fluent and consistent with the corresponding paragraph embeddings, it is essential to achieve a certain level of local smoothness in the paragraph embeddings space. Ideally, a slight variation in the input vector x would not cause a significant change in the resulting encoded vector $\mathcal{E}(x)$. Similarly, a small perturbation in the latent vector z should not lead to a significant change in the decoded vector $\mathcal{D}(z)$. Otherwise, the error accumulated in the diffusion process when generating z could result in an inaccurate realization of the desired semantics. To accomplish this, the denoising autoencoder is trained by substituting (Sub) input tokens with random tokens with probability p . The local smoothness is measured using the BLEU score ($\text{BLEU}_{\text{robust}}$) between the input x and the denoising output from corrupted input $\mathcal{D}(\mathcal{E}(\tilde{x}))$, where $\tilde{x} = \text{Sub}(x, p = 0.3)$. The level of injected noise will affect both the conversion error and the local smoothness, and it is important to strike a balance between the two.

Distributional Smoothness The diffusion model may face difficulty in learning a distribution, $p(z)$, that is highly multimodal, or has density that are associated with a large Lipchitz constant (*i.e.*, has abrupt changes). Therefore, we employ a text VAE (Bowman et al., 2016; Li et al., 2020) to encourage the posterior to take on a form close to a Gaussian distribution. Specifically, we parameterize $q(z|x)$ to be $\mathcal{N}(\mathcal{E}_\mu, \mathcal{E}_\nu)$ and maximize the objective $\mathcal{L}(\mathcal{E}, \mathcal{D}; x) = \mathbb{E}_{q(z|x)}[\log p(x|z)] - \beta \cdot \text{KL}(q(z|x)||p(z))$. Here \mathcal{E}_μ and \mathcal{E}_ν represent the posterior mean and variance predictions of the encoder \mathcal{E} , while the

hyperparameter β controls the strength of regularization. It is typically set to a small value to alleviate the notorious posterior collapsing issue (Bowman et al., 2016) in text VAE. To gauge the distributional smoothness of the paragraph embedding space, we select two examples, \mathbf{x} and \mathbf{x}' at random from the training set and interpolate their embeddings to compute $\mathbf{z}_{\text{INT}} = \frac{1}{2}\mathcal{E}(\mathbf{x}) + \frac{1}{2}\mathcal{E}(\mathbf{x}')$. We then evaluate the perplexity (PPL_{int}) of the decoded interpolation $\mathcal{D}(\mathbf{z}_{\text{INT}})$ using a GPT-2 model.

4.2 Planning then Decoding: A Latent Diffusion Model for Paragraph Embeddings

Training phase We now use the learned mean paragraph embeddings $\mathbf{z} = \mathcal{E}_\mu(\mathbf{x})$ to train a continuous-time latent diffusion model as in Fig. 2 while keeping \mathcal{E} and \mathcal{D} frozen. We conducted experiments using two types of conditioning signal: (i) class labels, such as positive or negative sentiment labels, and (ii) raw text, such as preceding context or the document to be summarized. For class labels, we learned a label embedding $\mathbf{y} \in \mathbb{R}^h$ to represent each class. For the raw text, we applied a conditional feature encoder τ to the input and used the hidden states at the last layer as $\mathbf{y} \in \mathbb{R}^{c \times h}$, where c represents the number of feature embeddings.

During training, we gradually add noise to \mathbf{z} via a cosine scheduler (Ho et al., 2020), and use a signal prediction scheme as the training objective (Kingma et al., 2021). For our denoising backbone model $\mathbf{F}(\cdot)$, we use a transformer block similar to the one in the DiT (Peebles & Xie, 2022) model. Specifically, we fed \mathbf{y} and the time embedding $\mathbf{t} \in \mathbb{R}^h$ into the model through two channels, namely cross-attention and adaptive layer norm (adaLN) (Peebles & Xie, 2022). For the cross-attention, the conditional embeddings \mathbf{t} and \mathbf{y} are concatenated into a sequence of length $c + 1$. The transformer block is modified to enable multi-head cross-attention to the conditional embeddings.

For the adaLN, we flattened and projected \mathbf{y} to a vector of \mathbb{R}^h using a linear projection layer. We then added the projected \mathbf{y} to the time embedding \mathbf{t} . Instead of directly learning dimension-wise scale and shift parameters (β and γ) in the standard Layer Norm (LN), these parameters are regressed from the sum of the embeddings. In addition, dimension-wise scaling parameters α are regressed and applied immediately before any residual connections within the transformer block. This has been shown to be efficient and effective in image diffusion models (Peebles & Xie, 2022).

Inference phase During the inference process, we start with random Gaussian embeddings and use a fixed number of steps T to generate the final \mathbf{z} . The resulting embeddings are then used as inputs for \mathcal{D} to generate the text using a *deterministic decoding* method like greedy decoding¹.

We applied the classifier-free guidance (CFG) (Ho & Salimans, 2021) during the inference steps. After each step, we apply a *dynamic thresholding* technique that was introduced in Imagen (Saharia et al., 2022) for post-processing. However, we do not use the rescaling step in Imagen because rescaling the paragraph embedding can completely alter the underlying semantics.

5 Experimental Setups

We tested the effectiveness of our model in three different conditional generation tasks including sentiment-guided generation, long-form text completion, and summarization. The tasks can require generating text of hundreds of tokens in length, making them suitable to assess model performance.

Datasets For the Sentiment-guided generation task, we used the TripAdvisor dataset provided by (Li et al., 2014). By exclusively selecting reviews with a rating of 1 or 5 and balancing the two ratings via subsampling, we acquired 218,509 reviews. For the text completion task, our model was assessed on two datasets: 1) the aforementioned TripAdvisor review dataset with postprocessing to remove reviews that are less than 20 or more than 512 tokens, result in 690,862 samples, and 2) one-tenth of the overall C4 datasets (Raffel et al., 2020), which contains 36.5M samples. For each sample, we extracted the initial two sentences from a paragraph as source context, and predicted the remainder of the text as target. The datasets were partitioned into training, validation, and test in the ratios of (0.96, 0.02, 0.02). For the summarization task, we use CNN/DailyMail (Hermann et al., 2015) and XS (Narayan et al., 2018). For fast evaluation of baselines, the evaluation is performed on 256 random subsamples from the test set.

¹The aim of \mathcal{D} is to accurately convert the \mathbf{z} into a meaningful text, thus deterministic decoding is desirable.

Automatic Evaluation Metrics Following previous work (Gong et al., 2023), we assess the **fluency** of the generation by computing the perplexity (**PPL**) under a GPT-2 large model. We use **Ent-n** (Zhang et al., 2018) and **DIST-n** (Li et al., 2016a) and self-BLEU (**S-BL**) (Zhu et al., 2018) to evaluate lexical diversity. We present DIST-n and Ent-n metrics solely at $n = 1$ owing to their strong correlation of the varying n values. We use **REP-n** to assess the extent of repetition in the generation following previous work (Welleck et al., 2019; Xu et al., 2022). For relevancy we use standard metrics following (Gong et al., 2023), including SacreBLEU (**BL**) (Post, 2018), ROUGE-L (**R-L**) (Lin, 2004) and BERTScore (**Score**) (Zhang et al., 2019). Details are provided in App. E.

AuBLEU: Evaluating Denoising Capability Our proposed model is a latent diffusion model, which differs from text diffusion models that operate directly on text or text embedding space. To comparing the denoising ability across different text diffusion models, we introduce a novel metric, named AuBLEU (**AuBL**). To compute the AuBLEU score, we first add varying levels of noise to each input text x by performing diffusion at T different time steps $t_0 < t_1 < \dots < t_T$, corresponding to a series of SNR $\omega_{t_0} > \dots > \omega_{t_T}$. Next, we pass each corrupted input under different ω to the denoising backbone model and obtain the predicted output $\hat{x}_i = F(x_{t_i})$. We then compute the BLEU score between each (\hat{x}_i, x) pairs and plot a curve with the x-axis representing $\alpha^2 = \frac{\omega}{1+\omega}$, where α^2 is monotonically increasing with ω and ranges from $(0, 1)$, and the y-axis indicating the corresponding BLEU score. Finally, we compute the area under curve to obtain the AuBLEU score (see App. C for more details).

Model Setups We used the BERT-large and GPT-medium models as initialization for the encoder \mathcal{E} and decoder \mathcal{D} respectively. The embedding dimension h was 1024, and the number of paragraph embeddings k was set to 16, as increasing the number did not result in significant improvement in performance. The learning rate was set to $2e - 4$, and β was set to $5e - 6$. For the latent diffusion model, the channel size was set to 1024 to match the embedding dimension h , and the number of heads was set to 16 with 28 transformer layers. The total size of the latent diffusion model was 533M. The feature encoder τ was also jointly learned, and was initialized with a T5-large encoder. We use DDIM throughout our experiments as it shows better performance than DDPM. In all our experiments, we use 30 diffusion steps to generate the final z , which strikes a good balance among the efficiency, diversity and relevance. In comparison, Diff-LM (Li et al., 2022) and Genie (Lin et al., 2022) report to use 200 steps and 2000 steps respectively to generate high-quality text. We set the CFG weights to be 2 and 5 for text completion and summarization tasks, respectively, based on generation performance on validation set. For summarization task, we also incorporate a shift noise scheduler based on (Hoogeboom et al., 2023). More details, including ablations on DDPM, number of diffusion steps and noise scheduler, are provided in App. E.

Baselines We compare our method with several baseline methods trained under Teacher Forcing scheme, including decoder-only Autoregressive LM finetuned on GPT-2 (**FT**), encoder-decoder (**Enc-Dec**) transformer model, and Variational Information Bottleneck (**VIB**) (Alemi et al., 2016). We initialized the FT model using GPT-2 large (774M), whereas encoder and decoder in the Enc-Dec/VIB models (695M/697M) are initialized with bert-large and GPT-medium, respectively. All the considered models are finetuned on the target datasets. We follow Li et al. (2022) to report the FT baselines with two decoding strategies, top-p sampling ($K=50, p=0.92$) and beam search (beam width 4), denoted as FT-sample and FT-search. We use top-p sampling for Enc-Dec/VIB generation. For summarization tasks, we finetune a T5-large model (770M) on the target datasets as baselines. We also compared two text diffusion models Diff-LM and Genie using their suggested configuration from the official repository. More details are in App. E.




Figure 3: Impact of the proportion of injected noise for learning Paragraph Embeddings on XSum dataset. Large substitution noise results in worse **BLEU_{clean}** but better **BLEU_{robust}** and **PPL_{int}**.

6 Results

6.1 Paragraph Representation Learning

It is essential to learn a paragraph embedding space that is both *accurate* and *smooth*. To this end, we examined the effect of various substitution probabilities p to the input tokens \mathbf{x} . Our findings, presented in Fig. 3, reveal that a smaller p results in a lower conversion error, as indicated by a higher reconstruction BLEU ($\text{BLEU}_{\text{clean}}$), albeit at the expense of local smoothness ($\text{BLEU}_{\text{robust}}$) and distributional smoothness (PPL_{int}). We performed a grid search of p with 0.1 increment based on 512 samples. Empirically, we observed that a weighted score $S_{\text{overall}} = 0.5\text{BLEU}_{\text{clean}} + 0.8\text{BLEU}_{\text{robust}} - 0.3\text{PPL}_{\text{int}}$ correlate well with downstream diffusion performance, leading to fluent and accurate generation for PLANNER (see App. A for more details). We finally opted for $p = 0.3$ for most datasets ², which strike a balance between conversion error and smoothness.

It is worth noting that there is an inevitable conversion loss, indicated by the fact that the $\text{BLEU}_{\text{clean}}$ is between 77 ~ 87 when generating hundreds of words (App. C). We observed that most of the lexical mismatch still maintain the similar semantics, with the exception of some name entity swapping. We show some paragraph reconstruction and denoising examples from our paragraph embedder in the App. A. We also include examples of interpolated generation from random paragraph pairs in the App. A. In general the transition of semantics is natural, indicating reasonable distributional smoothness of the paragraph embedding space.

6.2 Sentiment-Guided Generation

For sentiment-guided generation experiments, following previous works (Li et al., 2022; Hu et al., 2017; Keskar et al., 2019), we use a trained classifier to assess if the given sentiment is well-controlled in the generation. The trained classifier is initialized with BERT-

Arch.	PPL	ACC \uparrow	DIST/ENT \uparrow	S-BL \downarrow	Rep-4 \downarrow	Len
FT-sample	20.86	70.2%	0.13/6.154	0.96	5.86%	113
Diff-LM	101.97	83.6%	0.15/5.115	4.05	6.23%	66.2
Ours	51.12	94.9%	0.16/6.360	0.77	2.37%	161
Human	47.94	96.7%	0.17/6.737	0.48	2.17%	157

Table 1: PLANNER achieves high success rate (ACC) and diversity with less repetition in generating hotel reviews conditioned on sentiment.

large and finetuned on the training set, which yields an accuracy of 96.75% on the held-out test set. The results are provided in in Tab. 1. PLANNER outperforms the baseline approaches in generating long reviews at higher levels of accuracy. Although PLANNER using a greedy decoding mode is at least comparable with FT with top-p sampling in terms of diversity, and has lower repetition as assessed by Rep-4 in generation.

The perplexity of the text generated by PLANNER is close to that of human-written text. We provide examples of the generated text in App. D. Interestingly, as shown in App. D, with the same random seed but different controlling sentiment, PLANNER generates text with similar contents but different sentiments, suggesting the diffusion model may be able to disentangle the semantic space to certain extent. Unlike the autoregressive generation, the nature of the diffusion model allows the model to “regret” and iteratively refine on the current generations. In App. B, we demonstrate how the generation evolves over multiple time steps in a coarse-to-fine manner in PLANNER.

6.3 Long-form Text Completion

We further evaluate our model on the long-form text completion tasks. For text diffusion baseline, we compared our method with Diff-LM (Li et al., 2022) on hotel review dataset. We could not perform a comparison on the C4 dataset with Diff-LM due to the significant amount (thousands) of GPU hours required to train Diff-LM adequately. A Diff-LM running time estimation is available in App. E. The results are provided in Tab. 2. FT-search performed poorly in this open-ended generation task as its generation exhibited high repetition levels, consistent with findings in previous research (Holtzman et al., 2019; Xu et al., 2022). Although our approach also employs a deterministic decoding method, we observed that it produces text with low Rep-4 metric, signifying that PLANNER is effective in

²except for CNNDM dataset where we use $p = 0.5$

Arch.	PPL	DIST/ENT↑	S-BL↓	Rep-4↓	BL↑	R-L↑	Score↑	Len	AuBL↑
<i>Hotel Review</i> dataset									
FT-search	1.87	0.03/4.865	3.50	86.60%	0.62	5.2	0.39	179.51	-
FT-sample	15.51	0.14/6.455	0.88	4.49%	0.78	6.8	0.53	164.50	-
Enc-Dec	33.82	0.18/6.379	0.57	3.25%	0.47	7.3	0.54	94.03	-
VIB	36.89	0.19/6.481	0.54	3.15%	0.45	7.1	0.54	86.11	-
Diff-LM	178.30	0.13/5.560	3.57	4.54%	0.84	8.8	0.43	175.10	26.16
PLANNER	47.36	0.17/6.602	0.52	1.55%	0.77	7.9	0.55	168.08	38.55
Human	47.60	0.20/7.023	0.60	1.46%	-	-	-	181.29	-
<i>C4 subset</i> dataset									
FT-search	1.927	0.07/6.245	0.14	79.54%	0.77	5.2	0.37	154.88	-
FT-sample	12.244	0.25/7.136	0.44	7.01%	1.59	5.9	0.47	122.55	-
Enc-Dec	23.095	0.24/7.077	0.16	2.27%	1.92	7.5	0.5	118.07	-
VIB	19.701	0.24/7.003	0.16	2.62%	1.86	6.8	0.49	113.34	-
PLANNER	61.768	0.28/7.352	0.12	1.67%	2.04	7.7	0.51	111.89	36.77
Human	59.783	0.44/7.381	0.12	1.12%	-	-	-	107.56	-

Table 2: PLANNER enhances the diversity of text generation and minimizes the occurrence of repetition in open-ended text completion tasks.

reducing repetition through holistic iterative refinement throughout the inference steps in the diffusion process. Comparing with Diff-LM and other baselines, PLANNER achieved better diversity scores while maintaining comparable relevance scores. We also observe higher AuBLEU of PLANNER comparing with Diff-LM, indicating a potentially higher overall denoising strength of PLANNER (See App. C for more details). Some examples of the generated text are available in the App. D. We also observed PLANNER exhibits robustness towards prompts that are either repetitive or ill-composed, where FT failed (Fig. 1, App. F).

We further performed pairwise human evaluation on 300 examples of hotel review generation from each system on our internal crowd-source annotation platform. Each pair of text being presented to 3 judges in random order. The judges ranked the pairs for relevance, informativeness and human-like properties using a 3-point Likert-like scale. Overall judge preferences are shown in Table 3. A moderate preference can be observed for PLANNER over FT and VIB, except for human-like between PLANNER and VIB. We also observe that judges still prefer human responses over system generations in this task. Further details, including the human evaluation template used and interrater agreement analysis, are provided in the App. G.

Metric	Methods	Win	Tie	Loss
Rel.	Ours vs. FT	48.2%	9.2%	42.6% *
	Ours vs. VIB	50.7%	10.0%	39.3% **
	Ours vs. Human	39.3%	9.3%	51.3% **
Inf.	Ours vs. FT	55.1%	5.7%	39.2% **
	Ours vs. VIB	48.7%	8.0%	43.3% *
	Ours vs. Human	37.7%	8.7%	53.7% **
Hum.	Ours vs. FT	51.5%	8.4%	40.1% **
	Ours vs. VIB	40.0%	19.3%	40.7%
	Ours vs. Human	34.3%	17.0%	48.7% **

Table 3: Human evaluation on Relevance (Rel.), Informativeness (Inf.), and Human-likeness (Hum.). Statistical significant results: ** $p < 0.001$, * $p < 0.01$.

6.4 Summarization

We further conducted evaluation on summarization and present the results in Tab. 4. Summarization is less open-ended than the text completion task, thus a deterministic decoding approach like T5-search can produce high-quality text. Our evaluation shows that in comparison with T5-sample and Genie (Lin et al., 2022), PLANNER exhibits comparable Rouge-L scores, while improves other metrics. PLANNER achieves higher AuBLEU than Genie (See App. C for more details).

Owing to the sampling nature of the diffusion model, PLANNER and Genie yielded lower Rouge-L scores in comparison with T5-search, with single summary. To align with Genie’s evaluation, we provide the results with 10 random runs in Tab. 4, where for each document it generates 10 summaries and the one with the highest Rouge score is used. However, we note that these summaries with

best Rouge-L cannot be predetermined without an oracle summary. Comparing with T5-search, PLANNER generates more diverse and less repetitive summaries. However, the improvement is less conspicuous comparing with the results observed in open-ended text completion tasks. We show some generations in the App. D (Tab. 10).

Notably, the generated content may occasionally include hallucinations or errors, especially for name entities and digits (App. D, Tab. 11). Such occurrences can be attributed to either the conversion errors in \mathcal{D} or errors during the generation of paragraph embeddings, and requires further investigation.

Arch.	PPL	DIST/ENT \uparrow	S-BL \downarrow	Rep-4 \downarrow	BL \uparrow	R-L \uparrow	Score \uparrow	Len	AuBL \uparrow
<i>CNN Dailymail</i> dataset									
T5-search	58.24	0.38/6.977	0.23	1.54%	8.52	34.62	0.65	45.3	-
T5-sample	67.55	0.38/7.074	0.18	1.17%	5.37	30.21	0.63	49.3	-
Genie	151.8	0.32/6.674	0.20	0.82%	4.14	31.76	0.60	54.2	29.79
Genie ⁽¹⁰⁾	151.3	0.32/6.678	0.20	0.70%	3.90	37.72	0.60	55.1	-
PLANNER	49.95	0.37/7.250	0.14	0.89%	6.68	31.07	0.61	52.8	43.22
PLANNER ⁽¹⁰⁾	49.28	0.37/7.236	0.15	0.81%	12.16	37.14	0.65	53.1	-
Human	54.44	0.37/7.089	0.16	0.73%	-	-	-	55.0	-
<i>XSum</i> dataset									
T5-search	29.82	0.42/6.314	0.38	3.48%	6.27	38.03	0.72	17.3	-
T5-sample	36.45	0.43/6.497	0.28	1.30%	3.61	30.82	0.68	20.1	-
Genie	170.1	0.38/6.162	0.34	1.69%	3.27	36.20	0.68	17.7	31.80
Genie ⁽¹⁰⁾	170.9	0.38/6.173	0.32	1.45%	3.07	44.08	0.68	17.8	-
PLANNER	56.04	0.44/6.565	0.27	1.15%	5.65	36.07	0.7	18.5	48.77
PLANNER ⁽¹⁰⁾	56.57	0.45/6.540	0.30	1.24%	12.08	44.02	0.73	17.4	-
Human	47.09	0.49/6.564	0.28	0.64%	-	-	-	15.2	-

Table 4: For summarization task, PLANNER outperform Genie (Lin et al., 2022) in generation diversity and fluency while maintaining comparable Rouge-L scores. ⁽¹⁰⁾ indicates the maximum results after 10 runs, following (Lin et al., 2022).

6.5 Analysis

Running time We conducted inference time benchmarks of each method on a single Nvidia A100. For the sentiment-guided generation task, the autoregressive baseline is 5x faster than our method as the generation for all methods can be batched. For all other tasks, the varying input lengths make direct batchification for the FT baseline not straightforward. In these scenarios, the latent diffusion over a fixed number of latent codes offers computational advantages over a naive decoding of the FT baseline as the latent codes in our method can be conveniently batched. For the hotel review completion task, the generation of 256 samples took 378 seconds to complete, including 83 seconds for decoding and 295 seconds for diffusion generation with 30 generation steps. The FT baseline took 1,693 seconds to complete 256 generations. However, we note that padding or creating batches with the same length can potentially reduce the FT inference time. The Diff-LM algorithm required 397 seconds to produce 256 samples using 200 generation steps, which is comparable to ours. On the CNN-DM summarization tasks, our method took 682 seconds to generate 256 summaries. Genie’s generation took 3,822 seconds to generate 256 results. A naive decoding of T5 took 1,137 seconds to complete. XSum gives similar inference running time benchmark to the results on CNN-DM.

Generations over diffusion steps In App. B we provide generation examples for both summarization and sentiment-guided generation over different diffusion steps, which progress in a coarse-to-fine manner. The generation from early time step tends to be less fluent and generic. As the time approaches 0, the generation becomes more detailed. We presented quantitative results characterizing the evolution of the metrics over generation steps in App. B, Fig. 5. It revealed a clear trend of improvement in the majority of the metrics as the generation proceeds. Notably, most hallucinations occur during the late phase when more details are being incorporated. The model may excessively emphasize certain aspects, resulting in the correct generation being altered to an erroneous one (App. B, Tab. 9).

7 Conclusion

We present a two-stage latent text diffusion model that uses an autoencoder to condense lengthy texts into a limited number of paragraph embeddings, and a continuous time diffusion model that learns the distribution of these embeddings. Our proposed model alleviates the issue of repetition and advances generation diversity across different tasks. Compared to text diffusion models that perform diffusion solely on token or token embedding space, our method generates fluent text with improved diversity and reduced repetition. There may be toxicity or fairness issues in the dataset we used that we have not been able to identify. There are several limitations that warrant further investigation. Our work relies on an autoregressive decoder for converting latent representation into coherent text. It is possible to explore the feasibility of non-autoregressive decoders to bolster efficiency while minimizing conversion errors and hallucination in the generation. Furthermore, the classifier-free guidance approach results in a discrepancy between training and inference data distribution when feeding to the diffusion backbone. It would be interesting to investigate a “calibration” strategy for the latent code to better fit the data distribution during training.

Acknowledgement

We thank Yinfei Yang, Barry Theobald, Zhe Gan, Edouard Grave, David Grangier, Tatiana Likhomanenko, Richard Bai and Ronan Collobert for their critical suggestions and helpful feedback throughout this project.

References

- Alexander A Alemi, Ian Fischer, Joshua V Dillon, and Kevin Murphy. Deep variational information bottleneck. In *ICLR*, 2016.
- Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. *NeurIPS*, 34:17981–17993, 2021.
- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. Scheduled sampling for sequence prediction with recurrent neural networks. In *NeurIPS*, volume 28, 2015.
- Samuel Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. Generating sentences from a continuous space. In *CONLL*, pp. 10–21, 2016.
- Ting Chen, Ruixiang Zhang, and Geoffrey Hinton. Analog bits: Generating discrete data using diffusion models with self-conditioning. *arXiv preprint arXiv:2208.04202*, 2022.
- Pim de Haan, Dinesh Jayaraman, and Sergey Levine. Causal confusion in imitation learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/947018640bf36a2bb609d3557a285329-Paper.pdf.
- Shansan Gong, Mukai Li, Jiangtao Feng, Zhiyong Wu, and LingPeng Kong. Diffuseq: Sequence to sequence text generation with diffusion models. In *ICLR*, 2023.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. Teaching machines to read and comprehend. *Advances in neural information processing systems*, 28, 2015.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, volume 33, 2020.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. In *ICLR*, 2019.
- Emiel Hoogeboom, Alexey A Gritsenko, Jasmijn Bastings, Ben Poole, Rianne van den Berg, and Tim Salimans. Autoregressive diffusion models. In *ICLR*, 2021.
- Emiel Hoogeboom, Jonathan Heek, and Tim Salimans. simple diffusion: End-to-end diffusion for high resolution images. *arXiv preprint arXiv:2301.11093*, 2023.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P Xing. Toward controlled generation of text. In *International conference on machine learning*, pp. 1587–1596. PMLR, 2017.
- Bowen Jing, Gabriele Corso, Renato Berlinghieri, and Tommi Jaakkola. Subspace diffusion generative models. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIII*, pp. 274–289. Springer, 2022.
- Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*, 2019.
- Yoon Kim, Sam Wiseman, and Alexander M Rush. A tutorial on deep latent variable models of natural language. *arXiv preprint arXiv:1812.06834*, 2018.
- Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *NeurIPS*, 34: 21696–21707, 2021.
- Alex M Lamb, Anirudh Goyal ALIAS PARTH GOYAL, Ying Zhang, Saizheng Zhang, Aaron C Courville, and Yoshua Bengio. Professor forcing: A new algorithm for training recurrent networks. *Advances in neural information processing systems*, 29, 2016.
- Chunyu Li, Xiang Gao, Yuan Li, Baolin Peng, Xiujun Li, Yizhe Zhang, and Jianfeng Gao. Optimus: Organizing sentences via pre-trained modeling of a latent space. In *EMNLP*, 2020.
- Jiwei Li, Myle Ott, Claire Cardie, and Eduard Hovy. Towards a general rule for identifying deceptive opinion spam. In *ACL*, pp. 1566–1576, 2014.

- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. A diversity-promoting objective function for neural conversation models. In *NAACL*, 2016a.
- Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. Deep reinforcement learning for dialogue generation. In *EMNLP*, pp. 1192–1202, 2016b.
- Xiang Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. Diffusion-lm improves controllable text generation. *NeurIPS*, 35:4328–4343, 2022.
- Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pp. 74–81, 2004.
- Zhenghao Lin, Yeyun Gong, Yelong Shen, Tong Wu, Zhihao Fan, Chen Lin, Weizhu Chen, and Nan Duan. Genie: Large scale pre-training for text generation with diffusion model. *arXiv preprint arXiv:2212.11685*, 2022.
- Justin Lovelace, Varsha Kishore, Chao Wan, Eliot Shekhtman, and Kilian Weinberger. Latent diffusion for language generation. *arXiv preprint arXiv:2212.09462*, 2022.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 1797–1807, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1206. URL <https://aclanthology.org/D18-1206>.
- Myle Ott, Michael Auli, David Grangier, and Marc’Aurelio Ranzato. Analyzing uncertainty in neural machine translation. In *ICML*, pp. 3956–3965. PMLR, 2018.
- William Peebles and Saining Xie. Scalable diffusion models with transformers. *arXiv preprint arXiv:2212.09748*, 2022.
- Matt Post. A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, October 2018.
- Konpat Preechakul, Nattanat Chatthee, Suttisak Wizadwongsa, and Supasorn Suwajanakorn. Diffusion autoencoders: Toward a meaningful and decodable representation. In *CVPR*, pp. 10619–10629, 2022.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*, 2020. URL <http://jmlr.org/papers/v21/20-074.html>.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pp. 10684–10695, 2022.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. In *NeurIPS*, volume 35, pp. 36479–36494, 2022.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021a.
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *NeurIPS*, 32, 2019.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *ICLR*, 2021b.
- Sean Welleck, Ilya Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. Neural text generation with unlikelihood training. *arXiv preprint arXiv:1908.04319*, 2019.
- Ronald J Williams and David Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1(2):270–280, 1989.
- Jin Xu, Xiaojiang Liu, Jianhao Yan, Deng Cai, Huayang Li, and Jian Li. Learning to break the loop: Analyzing and mitigating repetitions for neural text generation. In *NeurIPS*, 2022.
- Lantao Yu, Weinan Zhang, Jun Wang, and Yong Yu. Seqgan: Sequence generative adversarial nets with policy gradient. In *AAAI*, volume 31, 2017.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert. In *ICLR*, 2019.

- Yizhe Zhang, Zhe Gan, Kai Fan, Zhi Chen, Ricardo Henao, Dinghan Shen, and Lawrence Carin. Adversarial feature matching for text generation. In *ICML*, pp. 4006–4015. PMLR, 2017.
- Yizhe Zhang, Michel Galley, Jianfeng Gao, Zhe Gan, Xiujun Li, Chris Brockett, and Bill Dolan. Generating informative and diverse conversational responses via adversarial information maximization. In *NeurIPS*, 2018.
- Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. Texus: A benchmarking platform for text generation models. In *The 41st international ACM SIGIR conference on research & development in information retrieval*, pp. 1097–1100, 2018.
- Zixin Zhu, Yixuan Wei, Jianfeng Wang, Zhe Gan, Zheng Zhang, Le Wang, Gang Hua, Lijuan Wang, Zicheng Liu, and Han Hu. Exploring discrete diffusion models for image captioning. *arXiv preprint arXiv:2211.11694*, 2022.

Appendix

A Variational Paragraph Embedder

A.1 Selection of substitution rate p

We observed when the value of p is within $(0, 0.7)$, there exists a correlation between the $S_{overall}$ and the PPL of the generation obtained from training PLANNER on the corresponding z (Figure 4). Performing a grid search on each task using diffusion models is an expensive process. Thus, we opted to use the surrogate $S_{overall}$ to choose the optimal p during the training of the paragraph embedder. However, it has been observed that an increase in the value of p leads to a deviation between the two. This could be attributed to a higher conversion error that occurs when p is excessively large.

A.2 Reconstruction, denoising and interpolation examples

In Table 5, we present examples that demonstrate the adeptness of the trained Variational Paragraph Embedder in providing clean and denoised reconstructions. Additionally, we showcase interpolation results (Table 6, 7) derived from two random sentences in the hotel review dataset. The interpolated paragraph is usually coherent and incorporates inputs from both sentences, characterizing the distributional smoothness of the latent space.




Figure 4: Impact of the proportion of injected noise for learning Paragraph Embeddings on XSum dataset. PPL_{int} and the PPL of the generation obtained from training PLANNER on the corresponding z at different noise level.

B Generation from PLANNER across multiple time steps

We provide generation examples for both summarization and sentiment-guided generation in Table 8 and Table 9. In general, it has been observed that generations progress in a coarse-to-fine manner. The early time step, which is close to 1, tends to be less fluent and generic. As the time approaches 0, the generation becomes more detailed and specific. Nevertheless, it has been noted that hallucinations may occur during the phase when more details are being incorporated. Occasionally, the model may excessively emphasize certain aspects, resulting in the correct generation being transformed into an erroneous one (see Table 9, in last two steps the highlighted text regarding “forensic DNA” is hallucinated, while the previous generations are more reasonable). The causes of such errors are under investigation for future studies.

In addition, we have presented quantitative results illustrating the evolution of the metrics during CNN-DM summarization generation based on 256 samples. These metrics are plotted in Figure 5. Our analysis has revealed a clear trend of improvement in the majority of the metrics as the generation process advances.

C Denoising strength comparison

To conduct a comparative analysis of text diffusion models’ denoising ability, we plotted the BLEU score under different signal-to-noise ratios (SNR) as shown in Figure 6. We use 20 time steps with an increment of 0.05 from $t = 0$ to $t = 1$ to compute the AuBLEU. The results indicate that our model exhibits a more uniform distribution of denoising ability across various levels of SNR compared to baseline models that operate on the word embedding space, as our model shows a stronger denoising ability when SNR is relatively small. Overall, PLANNER model achieves a higher AuBLEU. Note that we suffer from a conversion error resulting in a lower BLEU when the $SNR \rightarrow \infty$, *i.e.* $\alpha^2 \rightarrow 1$.

Clean reconstruction (<i>hotel review</i>)	
Original text	complaints : after a nights stay, i asked ther maid to clean our room (empty the waste baskets & make the bed). i returned at midnight and the room was never attended to. after 36 hours at mandalay, still zero maid service. i complained to management, and finally a maid came and clean our room. the manager said our first night stay would be free. upon check - out, we did not receive a free night stay. our bill was never modified, nor did they for us. what a joke. beware of the room's bathroom! the fossit water reeked of rotten mildew. i have a feeling they recycle their water... the stench was unbearable. many of the restaurants do not have bathrooms. you must exit the restaurant and use the bathrooms in the casino. very inconvienent. our in - room phone indicated that there was a message. i checked and no voicemail. the phone still blinked and blinked and blinked and blinked. still no voicemail. why is it blinking " new message ", if we don't have any messages?!? grrrrrr. the walls in the room are paper - thin. between girls laughing, to someone watching their tv.
Reconstructed text	complaints: after two nights stay, i asked the maid to clean our room (empty the wastebasket & make the bed). i returned at midnight and the room was never attended to. after 36 hours at the airport , still no maid service. i complained to management, and finally a maid came and clean our room. the manager said our first night stay would be free. upon check-out, we did not receive a free night stay. our bill was never corrected , nor did they compensate us for what they did . what a joke. beware of the room's bathroom! the faucet water reeked of rotten mildew. i have a feeling they rewash their water...the stench was unbearable. many of the restaurants do not have bathrooms. you must exit the restaurant and use the bathrooms in the lobby . very inconvenient. our in-room phone indicated that there was a message. i still received no phone message . the phone rang and rang and rang and rang . still no voicemail. is it a new message?? why we don't have any messages, " rudely "?? hmmm. the walls in the room are paper-thin. between her laughing, to someone watching their tv.
Denosing reconstruction (<i>hotel review</i>), noise level 0.3	
Original text	*** check out the bathroom picture *** i was in nyc by myself to watch some friends participate in the us olympic marathon trials. i figured with my wife back in portland, i could ignore the reviews and tough it out for a week. on the first night, i had a group of people enter my room with a key given to them by the front desk. i went to the desk and asked why in the world that could happen, let alone twice... he had no answer. i went back to bed and an hour later, again... the next morning i was so excited to get out for a run to literally escape the carter. i enjoyed a great run throughout central park ; when i returned i found three suitcases in the entry of my room. the owners entered the room a minute after i did and asked when i would be vacating the room so that they could unpack. we went to the front desk and complained and they said they 'hoped' it wouldn t happen again. want to unwind with tv. good luck. want to infect your lungs with mold, you will have better luck. seriously, i still have a cough. this place is unsanitary and absolutely unsafe.
Corrupted text	** [unused697] check exams the bathroom picture ** slams i was in nyc mead myself yankee 2016 some scotch ruin in the outfielder olympicna trials. i figured my gin [unused586] in portlandaki paramilitary could ignore inspected locoodon tough itwarkeron a 250. on [unused425] first rc, presentation traces a tribes of competitive enter my room with a key given to joint by the front . i went hope the fontana celeste oval norte in the world that could happengai let alone nickelodeon... he politics no answer. hancock went back reformed stool sousa an hour serge, again... consisting next morning i was so excited to get out for a run toelia escaperopolis napoleon. i enjoyed ct tian run throughout [unused965] park washed when lacrosse returned i found three suitcase white in the entry adapting room. the owners secretary the skirmish aivating after i did rhone drill i would be syriancating the room so memorials neutron sewer bobby [unused530]. would went to cassette front desk range complained and they said strikers byrd hoped' consistency wouldn ivision happen asylum. want to unwind with tv. good luck. want vega inump your lungs with mold, you will have bettercion. seriously waterways afforded still have a cough. this place is unsantamary and absolutely unsafe.
Reconstructed text	***check out the bathroom picture*** i was in nyc with my husband and some friends staying in the hudson hotel in nyc. i figured that my husband and in-laws could ignore the fact that it was not in a hotel. on the first night, i had a couple of people enter my room with a key given to them by the front desk. i went to the front desk to ask why in the world that could happen and let alone the hotel. he said no problem. i went back to the room an hour later, again... the next morning i was so excited to get out for a run to the theater...i continued to enjoy the run across the street. when i returned i found three suitcases in the entry way of the room. the owners had the key a while after i did so, so i would be aware that the room so far away from the elevators. i went to the front desk and complained and they said that the room wouldn't happen again. i want to unwind with tv. good luck. want to in on your vacation with you, you will have better luck. seriously, i still have a stench. this place is unsanitary and absolutely not safe.

Table 5: Reconstruction examples for clean reconstruction where input is not corrupted and denoising reconstruction where input is corrupted with 30% substitution noise. The mismatched text in the clean reconstruction is in **red**.

Sent A	Great resort beautiful views from the room. This was the nicest stay we have ever had. It was our honeymoon and we checked out of the Hilton in Waikiki after 1 night. Turtle Bay was a great resort. Big pool, waterslide and many restaurants and a great bar too.
$\tau = 0.2$	Great resort. Beautiful views from the room. This was the nicest stay we have ever had. It was our honeymoon and we checked out of the Hilton in Waikiki after 1 night. *Turtle Bay* was as nice. Big pool, waterslide and many restaurants and a great beach!!
$\tau = 0.4$	Great resort. Beautiful views from the room. This was the nicest stay we have ever had. We were on *honeymoon* and we checked out of the resort in the morning. The pool was as beautiful., *big pool and waterslide.*
$\tau = 0.6$	Fabulous resort. Beautiful views. Charming and entertaimentive service. We felt we were in a real resort. Only let down by the *pool*. The beach *was very old* and *smelled like mildew, and damp*.
$\tau = 0.8$	Huge lobby with beautiful chandeliers and furnishings. Overnightic stay and I thought we were in for a real treat. A step down when it comes to the room. *The smell was very old and smelled like mildew and damp*. The linens were very comfortable.
Sent B	Gorgeous lobby with beautiful chandeliers and furnishings. Overnightic smell. I thought we were in for a real treat. Only let down was the room. *The smell was so old and smelled of mildew and damp*. The linens appeared to be stale from the humidity

Table 6: Interpolation of short paragraph from the paragraph embedding space. $\mathcal{D}(z_A \cdot (1 - \tau) + z_B \cdot \tau)$

Sent A	the hotel is located in a very good part of downtown. walking distance from the vancouver harbour and english bay. we only stayed there for one night since we were just roaming around bc. we did not have the chance to try the restaurant downstairs but it looks like a very good setup. the lobby is pretty small so after check out and waiting for a taxi there were barely any seats to wait at. there is no gym, sauna, etc however, there is a ymca and steve nash sport centre facility nearby. again the location was great. it is a block away from granville street and they always have lots going on. " we got to see finger eleven doing a public presentation " that was neat for anybody that knows finger eleven. what i did notice that did not like for sure it is how noisy traffic can be downtown and how thin the walls were. we could here people closing and opening doors in the hallway and the only way to neutralized this was by turning on the ac. bathrooms were decent. the room overall was very clean and i had two queen beds in the room and we still had room to walk around them. tv was a flat screen. it had a mini fridge and the internet signal strength was good as well.
$\tau = 0.2$	the hotel is located in a very good part of downtown vancouver. walking distance from the vancouver harbour and english bay. we only stayed there for one night since we were just roaming around bc. we did not have the chance to try the restaurant downstairs but it looks like a very nice setup. the lobby is pretty small so after check in and waiting for a taxi there weren't even any seats to wait at. there is no gym, sauna, etc. however, there is a pharmacy and a 24hr fitness centre nearby. again the location was great. it is a block away from granville street and they always have something going on. "we got to see a show in a public theatre" that was neat for sure that anybody who knows the show. what i did notice did not matter for that it is how noisy vancouver can be downtown and how thin the walls are. we could here people opening and closing doors in the hallway and the only way to get to sleep was by turning on the ac. bathrooms were decent. the room overall was very clean and i had two queen beds in the room and i still had room to walk around it. there was a mini fridge. it had a flat screen tv and the soundproofing was good as well.
$\tau = 0.4$	i did not stay in a very good part of vancouver. walking distance from the vancouver harbour and english bay. we stayed there for one night since we were only there around 1pm. we didn't have the chance to try it because it looks like it's a new complex. the lobby is pretty simple so after check in and waiting for a taxi there weren't even any seats to be at a traffic stop. there is no gym, sauna etc. however, there is a pharmacy and a 24hr fitness centre nearby. while the location was great, it was a block away from stanley park and they still had to do everything on the weekend. we got to see "the metropolis" a neat commercial building that was perfect for sure when that is what you are looking for. what i did notice was not that it is as noisy as you can imagine and how thin the walls are. we could here people opening and closing doors in our hallway and the only way to get to sleep was in the morning. bathroom was decent. the room overall was pretty spacious and i had 2 queen beds in my room and i still had room to walk around it. there was a nice tv. it had a refrigerator and the sound proofing was good as well.
$\tau = 0.6$	i did not use a timeshare, but paid \$59 rate for a 3 night stay. we requested to be in the older building (i think there are 2 units there) and were in the newest part of the complex. it doesn't look like it's new, but the design is pretty typical of most other timeshare properties. you can see in the pictures from any of the rooms. it's on a back road, so unfortunately i cannot imagine, so please ask for it. we had a nice kitchen facility though. our shower was great though. it had a leak when we were there and they couldn't do anything about the water. they put in a huge construction crew thing that was noisy until, after the work out. second, what was strange is that the wall is paper thin. if you can hear everything, you have neighbors. we had to knock on the doors and keep our neighbors to the same way to drown it in the middle of the night. parking was horrible. the parking garage is very tight and almost every couple of spots are in need if you have an suv. i guess it was a nice decor, comfortable and it would take improvement. the view out front was great, but the noise from the street was a problem.
$\tau = 0.8$	i did not use a timeshare, but paid a daily rate for a 3 night stay. we requested to be in the newer building (i think there are 2 towers) and were in the newest part of the property. it's clean because it's new, but the design is typical of most other new timeshare properties. you can see disney in the distance from our room. it's on a back road, so obviously you cannot find it unless you ask for detailed directions. nicely decorated. we had problems though. our shower was getting hot enough and they had to repair it while we were there and couldn't use the shower for 6 hours. they put in a huge noisy air conditioner. however, even after the air conditioner was fixed. second and worst problem is the wall is paper thin. if you have neighbors, you can hear everything. we had to knock on the wall to tell our neighbors to keep it down at 1 in the morning. it was horrible. also, the parking garage is tight. almost every spot is hard to get into if you have an suv. i guess is a nice comfort, decor and breakfast. but would take away the opportunity to be in the newer building and have a little more privacy.
Sent B	i did not use timeshare points, but paid a daily rate for a 3 night stay. we requested to be in the newest building (i think there are 2 built) and got in the newest part of the property. it's clean because it's new, but the design is typical of most other new timeshare properties. you can see disney in the distance from our room. it's on a new road, so older gps cannot find it unless you pay for map updates. nicely decorated. we had some problems though. our shower was leaking big time and they had to repair it while we were there and couldn't use the shower for 6 hours. they brought in a huge noisy construction type dryer. however, even after the repair water was leaking. second and worst problem is the wall are paper thin. if you have neighbors, you can hear everything. we had to knock on the wall to tell our neighbors to keep it down at 1 : 00 in the morning. it was horrible. also, the parking garage is tight. almost every spot is hard to get into if you have an suv. i guess is your main concern is nice decor, comfort and clean, this would be ideal. but take into consideration the noise it's not so good.

Table 7: Interpolation of long paragraph from the paragraph embedding space. $\mathcal{D}(z_A \cdot (1 - \tau) + z_B \cdot \tau)$

D Generation examples

We present several examples of generation for each task in Table 10. Furthermore, we provide examples of problematic cases where the generated output may contain hallucinations or errors, particularly in terms of named entities and digits in Table 11.

E Experimental setup

E.1 Metrics

We evaluate the generation of the model using automatic evaluation and human evaluation. Following previous work (Gong et al., 2023), we assess the **fluency** of the generation by computing the

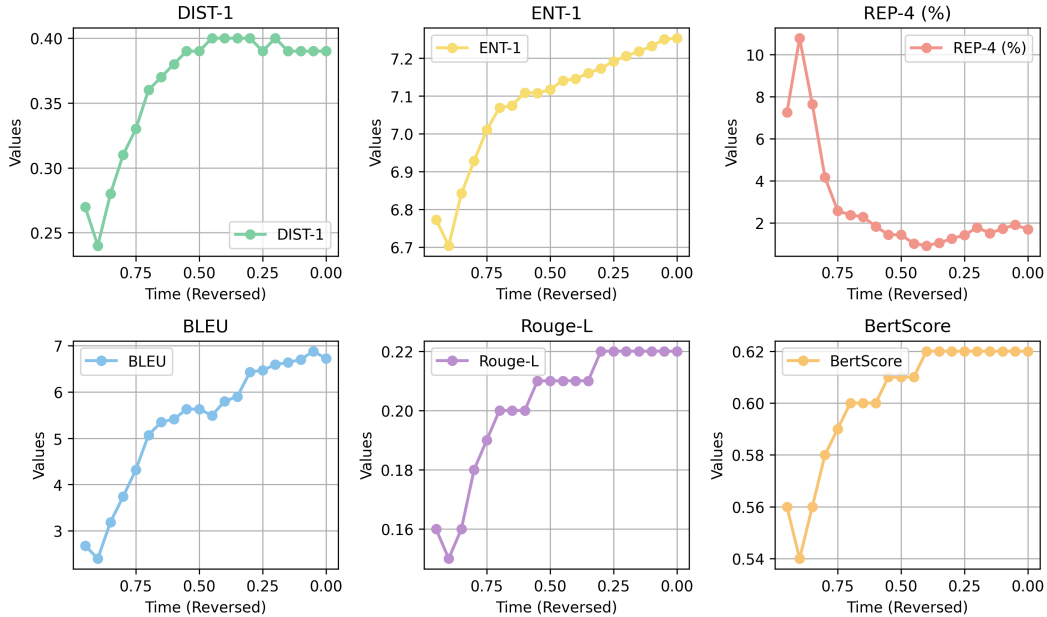


Figure 5: The evolution of the metrics during CNN-DM summarization generation from PLANNER based on 256 samples. We observe a clear trend of improvement in the majority of the metrics as the generation process advances.

Time 0.90	I the hotel is at this to beynd the staff here is great. I'll be back here soon. Highly recommend this place.
Time 0.80	Well stocked, great service at a great service. I've stayed at this beautiful hotel.
Time 0.70	I can't say enough. Great staff. Great meals. They have great service too. The food is tasty. Great place to bring the kids. Great pool. Great grounds. Great service.
Time 0.60	I can't say enough about this hotel. Great restaurants. They have great service. It's a very family friendly hotel. Glam crowd at the pools. Great golf courses. Two pools and two hot tubs. Great for kids.
Time 0.50	I can't say enough about this fabulous resort. They have fantastic service. It was a very hip atmosphere. The decor is cute and cushy. Two pools, two water slides and two hot tubs. Gorgeous grounds with a great golf course. Well done.
Time 0.40	I can't say enough about the Bahia. They have refreshingly funky. It was a very modern atmosphere. The decor is cozy and comfy. Loved the grounds, two swimming pools and two restaurants. Thank you for a great time. Looking forward to going back. Well done.
Time 0.30	I can't say enough about the Bahia. They have refreshingly funky. It was a very modern atmosphere. The decor is cozy and comfy. Loved the grounds, two swimming pools and two restaurants. Gotta go for the coffee. Thank you for a wonderful time. Thanks for a wonderful time.
Time 0.20	I can't say enough about the Bahia. They have refreshingly funky. It was a very modern atmosphere. The decor is cozy and comfy. Loved the two pools, two hot tubs and two restaurants. Thank you for a great time. Looking forward to coming back.
Time 0.10	I can't express enough about the bartender, Bahia. They have great service. It was a very modern atmosphere. The ambiance of the place is cute. The grounds are lush. The venue boasts two pools, two hot tubs. Great restaurants. Great bar. Great service. Great location. Thanks.
Time 0.00	I can't express enough about the bartender at this establishment. They have a modern and creative vibe. The ambiance of the place is simply adorable, with chic decor that adds to the overall experience. The venue boasts two restaurants and two hot tubs, which is quite impressive. Bravo to the bartender! Thanks.

Table 8: Generation from the diffusion model with 10 steps on hotel review dataset with positive sentiment. The generation progress in a coarse-to-fine manner.

perplexity (**PPL**) under a GPT-2 large model. For text completion, the context is concatenated with the generation, and only the generated text's perplexity is considered. We use **Ent-n** (Zhang et al., 2018) and **DIST-n** (Li et al., 2016a) and self-BLEU (**S-BL**) (Zhu et al., 2018) to evaluate lexical diversity. The Ent-n quantifies the entropy of the empirical n-gram distribution of the text generated, while the DIST-n metric calculates the proportion of unique n-grams among all n-grams. We present DIST-n and Ent-n metrics solely at $n = 1$ owing to their strong correlation despite the varying n values. The self-BLEU metric is used to compute the inter-example BLEU score,




Figure 6: The BLEU score under different SNR for PLANNER and baselines. The AUC under these curves are the AuBLEU metrics.

which evaluates cross-example diversity. We use **REP-n** to assess the extent of repetition in the generation following previous work (Welleck et al., 2019; Xu et al., 2022). The REP-n is defined as $1 - |\text{Unique n-gram}|/|\text{n-gram}|$.

E.2 Model setups

We used the BERT-large and GPT-medium models as initialization for the encoder \mathcal{E} and decoder \mathcal{D} respectively. The embedding dimension h was 1024, and the number of paragraph embeddings k was set to 16, as increasing the number did not result in significant improvement in performance. The learning rate was set to $2e-4$, and β was set to $5e-6$. During training, 30% of the input tokens were substituted to a random token.

For the latent diffusion model, the channel size was set to 1024 to match the embedding dimension h , and the number of heads was set to 16 with 28 transformer layers. The total size of the diffusion model was 533M. The feature encoder τ was also jointly learned, and was initialized with a T5-large encoder. For text completion and summarization tasks, we used the first 256 hidden states from the last layer as \mathbf{y} . We use DDIM throughout our experiments as it shows better performance than DDPM across the board. To enhance the summarization performance of the model, we incorporate a shift noise scheduler with noise_shift= 4 based on Hooeboom et al. (2023). This scheduler encourages the model to concentrate more on the high noise level phase of the training process.

Following (Ho & Salimans, 2021), we use a CFG dropout ratio of 10% during training. During inference, for text completion tasks, we set the CFG weights to be 2, while for summarization tasks, we set the CFG weights to be 5, based on performance on the validation set. We use greedy decoding across all the tasks in PLANNER to decode text from predicted latent embeddings as we do not seem noticeable improvement of performance by using a beam search decoding method.

We utilized 4 Nvidia A100 GPUs to train every model until convergence, based on validation loss. While training the paragraph embedder, the batchsize was fixed to 48. It took about 20 to 40 hours for each dataset to complete 20 epochs of training. The diffusion model was trained with a batchsize of 12, which lasted for 50 hours for summarization and text completion tasks. The training of sentiment-guided generation task only took approximately 20 hours until convergence. FP16 was employed throughout the duration of the training process for better efficiency.

E.3 Text diffusion baseline configurations

Our experimental setup for Diff-LM is based on Diff-LM’s official implementation and configuration described in Li et al. (2022). We follow Diff-LM to use employ BERT-base with 80M parameters as the backbone model and utilize a square-root noise schedule along with a diffusion forward step of 2000 and decoding steps of 200. Additionally, the embedding dimension for our models is set to 128. We use a sequence length of 128 for sentiment-conditioned generation and 256 for long-form text completion tasks.

As reported in Diff-LM, it requires approximately 5 hours to execute 200,000 iterations when trained on the E2E dataset. However, when training with the larger ROCStories dataset, which contains 98,000 five-sentence stories, it has been suggested that the algorithm be trained for at least 800,000

iterations, which requires over 20 hours of GPU time. Notably, the C4 subset contains 372.4 times more documents than ROCStories, even when the document size is not considered. As a result, at least 7,448 GPU hours would be required to adequately train the algorithm for 800,000 iterations using C4.

The official implementation of Diff-LM employs a fixed-length decoder that contains some special tokens, including paddings. As a result, it produces high Rep-4 scores. To ensure a more objective evaluation, we performed additional post-processing to eliminate paddings and recomputed the scores based on post-processed generation.

For Genie (Lin et al., 2022), we used their official implementation as well as their fine-tuned checkpoints for XSum and CNN/DailyMail datasets, which are released officially as per Lin et al. (2022). These checkpoints are optimized using a 6-layer transformer as the encoder, pre-trained on a large 160G corpus for 5 million steps. Furthermore, a 6-layer cross-attention transformer is employed for denoising. Additionally, the latent variable dimension is set to 768, while the embedding dimension is set to 128. Genie’s configuration also includes a uniform time schedule with 2,000 as the diffusion steps.

E.4 Ablations on DDPM, diffusion steps and noise scheduler

We present ablations on DDPM vs DDIM, and PLANNER using different diffusion steps and different noise scheduler in Tab. 12. DDIM is better than DDPM across all of our experiments in most of the metrics except a slight drop in terms of diversity. We observed that more steps will typically improve the diversity score at a cost of relevance and inference speed. We also compared the cosine scheduler with the beta linear scheduler (Rombach et al., 2022). The cosine scheduler worked better in our experiments. For the summarization tasks, we found that using a noise shift Hooigeboom et al. (2023) of 4 improves the Rouge-L by around 5%, comparing to a vanilla setting with noise shift of 1.

F Text completion with repetitive prompt

We present examples in Table 13 of generation with an ill-composed prompt for hotel review generation. The results reveal that the FT baselines tend to generate repetitive text. Although sampling mitigates this issue to some extent, self-reinforcement of repetition still occurs during generation. In contrast, our model exhibits significantly less repetitive generation.

G Human evaluation

We screen the judges using 10 random screening questions, the judges pass 80% can participate our evaluation. The interrater agreement assessed by Krippendorff’s alpha is 0.5958. The template used for human evaluation is provided in Figure 7

Guidelines

Instruction:
Imagine you're reading a hotel review. Compare the two candidate text completions (**Follow-up**) for a given partial review as **Context**. Using the radio buttons, determine which is more relevant/consistent and with specific details in light of the preceding context. Then decide which of the two Follow-up Text seems to be more human-like, again in light of the previous context.

There is room for subjectivity here, but you should primarily consider whether it would be reasonable to think that a person who read the review might find the follow-up text relevant/consistent and contentful. A generic or bland follow-up like "This hotel is awesome" may be relevant/consistent, but not especially contentful.

Some text are incomplete. Please make judgement on the partial text regardless of whether they are complete. Please ignore upper/lower case, punctuation, spacing, and other minor formatting issues. It is expected that you will need to spend at least 10 seconds working on this page.

For example, given a context:
Context: we spent one night here and very much enjoyed it. the room was very spacious, and we liked the unique kington decor.

Follow-up 1 (Not relevant/consistent): Same phone number 17 years later and a real live person!!! No recording!!! We booked a room for 3 days. We had a wonderful time! The staff is top notch, so courteous and helpful. We love the extra touches of coffee, tea, cookies, fruit, newspapers 24-hours 7 days a week.

This text is not consistent with the context. To assess its relevancy and consistency, you should determine if the follow-up text fits well with the context. Some usual problems include irrelevant content, inconsistent content, and inconsistent writing style.

Follow-up 2 (Good): I got a kick out of the animal-print bathrobes. We slept very well - the bed was very comfortable, the window shades blocked out most light, and we could not hear any noise from the hallway or adjoining rooms. Location is very good, as it is an easy 15 minute walk to the inner harbor area. The only major disappointment of our stay was the valet service.

Relevant follow-up.

Contentfulness: Which follow-up has more specific details?

Answer: Follow-up 2

Follow-up 1 (Not contentful): We really love this hotel and they have great staff! We will come again!

This text is bland and generic as it contains very few specific details. It's not interesting and informative.

Follow-up 2 (Good): I got a kick out of the animal-print bathrobes. We slept very well - the bed was very comfortable, the window shades blocked out most light, and we could not hear any noise from the hallway or adjoining rooms. Location is very good, as it is an easy 15 minute walk to the inner harbor area. The only major disappointment of our stay was the valet service.

A follow-up with many details

RELEVANCE/CONSISTENCY: Which follow-up is more relevant and consistent to the context?

Answer: Follow-up 2

Follow-up 1 (Not human-like): This hotel was the convenient spot to the park to shop. We loved the little park, although we loved a lot, we loved the outdoor deck. And much more. It was well staffed and well kept. I would stay here again. Highly recommend.

It is unlikely that this text was written by a human. Common issues that arise with non-human generated text include disfluency, lack of readability, and repetitive language. You may use your own judgement to determine if this text was composed by a human or not.

Follow-up 2 (Good): I got a kick out of the animal-print bathrobes. We slept very well - the bed was very comfortable, the window shades blocked out most light, and we could not hear any noise from the hallway or adjoining rooms. Location is very good, as it is an easy 15 minute walk to the inner harbor area. The only major disappointment of our stay was the valet service.

A follow-up which is likely written by a human

WHICH IS MORE HUMAN: Which follow-up seems more likely to have been made by a human than an AI model?

Answer: Follow-up 2

Task ID: 01G10W9F0VQK8CNC66J7LSJH

Show Rating: Select a rating... View Rating: 01G10W9F0VQK8CNC66J7LSJH-0 View Rating: 01G10W9F0VQK8CNC66J7LSJH-1

Context: the crystal pier is in the pacific beach section of san diego. if i remember correctly, this is one of five beach communities north of the center city.

Follow-up #1: The Crystal Pier is in the Pacific Beach section of San Diego, one of five beach communities located north of the city center. It is popular among locals and visitors alike due to the large selection of fun amenities, historic attractions and great restaurants, such as The Inn at the Open, the Buena Vista, and the Old Town Trolley. The city on the bay also offers plenty of entertainment and activities, making it a great destination for a vacation. Hotel Monaco is located nearby on Geary Street and provides clean, historic rooms with friendly service. Despite some reviews noting the need for maintenance, the hotel is a great choice for those looking for a place to stay in the area. The staff is friendly and the six-lane adjacent hotel garage allows for easy access from the street. The hotel rooms in San Francisco are usually better during this time of year, but the hotel offers a great atmosphere that can't be found elsewhere - with its mix of antiques, retro design, and club-like vibes.

Follow-up #2: is the home of one of the most influential members of the shibboleth family, Omor-Omor. He is a major player in the pacific beach community and also the one who first gave the shibboleth tribe its name, which is now called the shibboleths. The shibboleths are a tribe that uses the crystal pier as a weapon. They are believed to be the original barbarians, though they have been seen in other races like the aarakocra. The shibboleths are famous for killing or driving away enemies who were pursuing them out of fear for their safety. It is also likely that they have been traveling through this part of the world for thousands of years. The crystal pier is situated off the west coast of the pacific on the coast of south pacific. On this pier the shibboleth tribe resides. The shibboleths have the ability to generate waves and can create strong waves through their crystalline skin that can create tremendous amounts of damage in rapid succession. This power gives the

RELEVANCE/CONSISTENCY: Which follow-up is more relevant and consistent to the context?

Follow-up #1

Follow-up #2

Equally

Contentfulness: Which follow-up has more specific details?

Follow-up #1

Follow-up #2

Equally

WHICH IS MORE HUMAN: Which follow-up seems more likely to have been made by a human than an AI model?

Follow-up #1

Follow-up #2

Equally

[Optional] Any additional comment?

Figure 7: Template for human evaluation.

Document	Washington (CNN)Maryland authorities said Wednesday that a former state correctional officer has been arrested in connection with a recent spate of shootings, including one on the Intercounty Connector in Maryland and one at Fort Meade, where the National Security Agency office is located. Officers stopped Hong Young, 35, of Beltsville, Maryland, at around 10:30 p.m. Tuesday. The officers recognized Hong's vehicle – a 1999 Lincoln Town Car – as matching authorities' description of a car seen in surveillance footage near some of the shootings. A gun in the car matched evidence found at the shootings , authorities said at a press conference, and Young was arrested. Young is in the hospital and under police guard, though when reporters asked why he was being treated, officials would only say he was arrested without incident. He is charged with attempted first-degree murder, first- and second-degree assault, weapons violations and reckless endangerment. Young worked as a correctional officer at a Jessup facility from 2012 until his resignation in 2014, Maryland Secretary of Public Safety Stephen Moyer said. There was nothing significant in his employee file, Moyer said. Police said that there are no links to terrorism, and no motive has been determined. No one was killed in the five shooting incidents, four of which occurred Monday and Tuesday, according to police reports. – February 24 in Hanover, Maryland, a man who had stopped at a Costco said a man pulled up beside him in a Lincoln Town Car at 7:30 a.m. and began firing at him. The victim's vehicle was hit several times and the victim was grazed. The assailant drove away. – March 2 in Laurel, Maryland. Police received a call at 2:50 a.m. that shots had been fired at a Walmart. There were no damages or injuries. – March 2 in Columbia, Maryland. A call came in to law enforcement at 4:51 a.m. about shots fired at a movie theater at Columbia Mall. Surveillance footage captured a Lincoln Town Car at about the same time shots were fired, police said. Though several employees were there, no one was hurt, authorities said. There were bullet holes in the theater glass and a shell casing was found at the scene. – March 3 in Prince George's County. Multiple shots were fired at an overpass on the InterCounty Connector in the afternoon, striking a tree service truck with two passengers inside. – March 3 at Fort Meade. Shots struck a building near the NSA office at about 6 p.m. Along with the gun, evidence shows Young was the shooter in all but the Walmart incident, though that investigation is continuing, police said. Though no one was killed in the incidents, they stirred memories of the deadly Washington, D.C.-area sniper attacks in 2002. Ten people were killed in Washington, Maryland and Virginia during that rampage, which went on for three weeks. CNN's Holly Yan and Laurie Ure contributed to this report.
Time 0.95	Police: man in "gunman" shooting in police shooting, police say. Police in a car, police in a vehicle, police say.
Time 0.90	Police: man are in police shooting, murder in police shooting, police say. Police in a vehicle found in police shooting, police say.
Time 0.85	Police: man arrested in one murder, one in shooting, police say. Police find a car, a police vehicle in the suspect, police say.
Time 0.80	Police: Young man arrested in attempted murder, one in shooting, police say. Officers in a car matching the vehicle found in surveillance footage, police say.
Time 0.75	Police arrest man, man charged with attempted murder, shootings, police say. Officers say a car matching police description found in surveillance images, police reports.
Time 0.70	Police arrested Lee Young, with attempted murder, assault charges, weapons violations, authorities say. Officers found a car matching surveillance footage found in the vehicles, reports say. One of the shootings, no one was in the car.
Time 0.65	Police arrested Hong Lee, 35, with attempted murder shootings, weapons violations, authorities say. Vehicle in his car matched surveillance footage found in surveillance images, police say. Only one in five shootings were no one was killed in Maryland, police say.
Time 0.60	Police arrested Hong Young Lee, 35, with attempted first-degree murder, authorities say. Car in the car matched surveillance footage found in surveillance images, police say. One of the five shootings occurred in the home. No one was killed in Maryland, but no motive has been determined.
Time 0.55	Police arrested Hong Young Lee, 35, with attempted first-degree murder, authorities say. Car in the car matched the surveillance images found in surveillance footage, police say. No one was killed in the five shootings.
Time 0.50	Police arrested Hong Young, 35, for attempted first-degree murder, authorities say. Car in the car matched the evidence found in surveillance footage, police say.
Time 0.45	Police arrested Hong Young, 35, for attempted first-degree murder, authorities say. Car in the car matched the evidence found in surveillance footage, police say.
Time 0.40	Police arrest Hong Young, 35, in attempted first-degree murder, authorities say. A car matched surveillance images found in surveillance footage.
Time 0.35	Police arrested Hong Young Li, 35, in attempted first-degree murder, authorities say. A car in the car matched surveillance images found in surveillance footage.
Time 0.30	Police arrested Hong Young Li, with two attempted shootings, assault charges, authorities say. A gun in the car matched surveillance images found in surveillance footage.
Time 0.25	Police arrested Hong Young, with two attempted shootings, assault charges, authorities say. A gun in the car matched surveillance images found in surveillance footage.
Time 0.20	Police arrested Hong Young, with two attempted shootings, assault, authorities say. A gun in the car matched surveillance images found in surveillance footage.
Time 0.15	Police arrested Hong Young, 35, with attempted first-degree murder, assault, authorities say. A gun in the car matched surveillance images found in surveillance footage, police say. No one was killed in the five shootings. No motive has been determined.
Time 0.10	Police arrested Hong Young with attempted first-degree murder, assault, authorities say. A gun in the car matched forensic identification found in surveillance footage, police say. No one was killed in the five shootings. Weapons violation. No motive has been determined.
Time 0.05	Police arrested Hong Young with attempted first-degree murder, assault, authorities say. A gun in the car matched forensic DNA found in surveillance footage, court documents show. No one was killed in the five shootings. Weapons violations. No motive has been determined.
Time 0.00	Police arrested Hong Young after attempted shootings, assault, authorities say. A gun in the car matched forensic DNA found in surveillance footage, police say. No one was killed in the shootings in Maryland. Weapons violations, police say. No motive has been determined.

Table 9: Generation from the diffusion model with 20 steps on CNNDM dataset. In last two steps, the highlighted text regarding “forensic DNA” is hallucinated, while the previous generations (e.g. “A gun in the car matched surveillance images”) are more reasonable.

Semantic generation (<i>hotel review</i>)	
PLANNER (Negative)	I've received several complaints about the amenities of hotels, although the rooms are standard La Quinta, and the price is a bit low. I booked a room for a family of four. Given the low price, I expected to tolerate the motel's unappealing colors, but that's probably par for the course. I read the reviews and chose this hotel for a night sleep, but I was incredibly disappointed, to say the least. Part of the issue might have been the noise from the freeway, compounded by the fact that I could hear the AC from my neighbor's room. There was also a noise that sounded like "flooring". The worst part was that the bed sheets had no cover, and the AC draft was felt throughout the room. The room itself was cramped, slightly outdated, and uncomfortable. The furniture was old, but there wasn't much else in terms of extras. If it weren't for the free drinks and the front desk service, I would have rated this hotel much lower than most others.
PLANNER (Positive)	I've got numerous praises about the amenities of the hotel. Rooms are typical of La Quinta's excellent standards, and the price is low surprisingly. I booked a room from they website. I was surprised for the hotel's distinctive color scheme, quite a signature style. This hotel I selected for overnight stay, and I was overwhelmingly satisfied with no complaints. A key element of the charm was the profound tranquility throughout the hotel. Despite being conveniently near the freeway, the excellent soundproofing ensures a peaceful stay. Which was great was that the bed sheets were so soft and comfortable, and the gentle breeze from the AC added to the overall soothing ambiance. The room itself was cozy, with a charming vintage vibe, and supremely comfortable. The furniture was classic, giving a nostalgic touch. The complimentary drinks and the outstanding service at the front desk were delightful, making this hotel a great choice.
Retrieved from training set	The hotel is located in a great area, close to a lot of things. After reading a lot of the previous reviews I booked it and decided to see for myself. The checkin was a little frantic but there were a lot of people so I wouldn't count that against them. The gentleman that checked me in was extremely polite and apologized several times for the wait, which hadn't been that long. I went up to my room on the second floor and was surprised that it was bigger than I thought it would be. I did have the tiny bathroom that previous posters commented on but it was fine. My only complaint would be the noise level. I was in a room that faced the street and the noise was unbelievably loud. If I stayed here again I would ask for a room which doesn't face the street. It was nice to have a Starbucks right off the lobby, as well as a coffee maker in the room. So, for the money, I would recommend the Hotel Diva.
Text completion (<i>C4 subset</i>)	
Context	Operating a successful dental practice isn't easy. That's why we've used our industry knowledge and expertise in practice management to create an eBook that offers tips on how you can create and maintain a financially healthy and profitable dental operation.
PLANNER	This e-book is akin to a popular one, if not more stylish. It provides all the benefits you need for your accounting endeavors for your dental practice. Unlike other books, this practice management guide will enlighten you on what it's like to pay bills and kickstart your business. You'll learn how to effectively stage your practice management, and the book offers valuable tips and financial ideas for launching your dental business. It delves into a fascinating topic, presenting the steps you need to enhance your financial productivity and support your financial growth plan. You'll discover how to manage your ongoing fiscal and accounting needs, and receive tips that will enable you to achieve timely results and make informed financial decisions. The book includes quick and detailed insights on driving numbers and finding solutions. Learn how to start your accounting practice and achieve the ultimate goal of creating a profitable and successful business. We appreciate your feedback! Leave a review, and we'll send a book for you. We offer examples of the key things you should consider when deciding to start your dental business.
Summarization (<i>CNN-Dailymail</i>)	
Document	(CNN)Actress Liza Minnelli has entered a rehab facility for her addictions, a spokesman said. "Liza Minnelli has valiantly battled substance abuse over the years and whenever she has needed to seek treatment she has done so," said spokesman Scott Gorenstein. "She is currently making excellent progress at an undisclosed facility." The 68-year-old has struggled with addictions to alcohol and painkillers in the past. Minnelli won an Oscar in 1973 for her performance in "Cabaret." CNN's Topher Gauk-Roger contributed to this report .
PLANNER	The actress Liza Minnelli is "making excellent progress" at a undisclosed facility, spokesman says. She won an Oscar in 1973 in "Cabaret"
Summarization (<i>XSum</i>)	
Document	Japanese exports rose by just 0.6% in September from a year earlier, much slower than in previous months. The data is seen as reflecting a slowdown in Japan's top trading partner, China. However, the Nikkei 225 closed 1.9% higher at 18,554.28 on hopes the weak data would trigger more stimulus efforts from the government. In China, the mainland's benchmark Shanghai Composite closed down 3.06% at 3,320.68. Hong Kong's Hang Seng remained closed on Wednesday for a public holiday. In Australia, the S&P/ASX 200 rose by 0.2%, closing at 5,248.30. Shares in mining giant BHP Billiton rose 1% after it maintained its full-year production guidance for iron ore. Slowing growth in China has been weighing on mining stocks in recent weeks as the country is the main importer of Australian commodities. South Korea's benchmark Kospi index rose marginally by 0.2% to finish at 2,042.98. Shares of LG Electronics jumped by more than 14%, boosted by a fresh supply agreement with General Motors.
PLANNER	Asian markets closed higher after weak economic data indicated slow growth due to lower exports.

Table 10: Generated examples.

Summarization (<i>CNN-Dailymail</i>), hallucination issue	
Document	(CNN)A Los Angeles police officer charged with murder in the shooting death of a local truck driver had his father's help in fleeing California and traveling to Texas, court documents filed Thursday said. Henry Solis, a Los Angeles Police Department rookie, is accused of killing 23-year-old Salome Rodriguez Jr. early in the morning of March 13 in Pomona. Solis disappeared shortly after the shooting. Items left by Solis at the crime scene allowed Pomona police detectives to identify him "as the individual responsible for killing Rodriguez," according to an affidavit signed Thursday by FBI Special Agent Scott Garriola. The document goes on to say that Solis made "incriminating statements about his role in the murder" to friends, family members and others, and said he would never be seen again. The affidavit says his father, Victor Solis, quickly left his home in Lancaster after receiving a call from him. The elder Solis was seen a day later in El Paso, Texas, at the home of family members. He told FBI agents he drove his son to El Paso, dropped him off at a bus station and doesn't know where he is. Federal authorities haven't said if Victor Solis would face charges for helping his son leave California. CNN's Tina Burnside contributed to this report.
PLANNER	Henry Solis is accused of killing Delara Salome Guzman Jr. , Solis disappeared shortly after the shooting. Victor Solis helped his son leave California.
Text completion (<i>C4</i>), fluency issue	
Document	Two men have died after being crushed by timber that fell on top of them inside a shipping container. Authorities were called to an industrial site in Wetherill Park , west of Sydney, shortly after 11am on Thursday where they worked to release three men. Police found the three men inside a shipping container on a property on Victoria Street. Authorities were called to an industrial site in Wetherill Park, west of Sydney, shortly after 11am on Thursday where they worked to release three men . They had been unloading or loading timber when it fell on them. Two men died at the scene as a result and the surviving man was taken to nearby Fairfield Hospital in a serious condition. Police say the shelving in the shipping container moved and trapped the men. Police and WorkCover NSW are investigating the incident and a report will be prepared for the Coroner. Two men died at the scene as a result and the surviving man was taken to nearby Fairfield Hospital in a serious condition .
PLANNER	Authorities were called to Worsley Park , west of Sydney, Sydney. Police found three men inside a shipping container after 11am on Thursday. Two men died at the scene as a surviving man was taken to nearby hospital. NSW police are investigating the incident.
Text completion (<i>C4</i>), fluency issue	
Prefix	Can't quite find what you're looking for? Try our handy site search.
PLANNER	Fill our search engine form and try us and we'll try to get back to your desk on time. We have a reliable staff that provides professional service, whatever you require - not just a specific search function. But we also let's you enjoy on-line selection of different products, storage, protection, tools, books and accessories, and many more products. Simply what you need! Our main site is the updated web site site you can visit. We don't store any coupons, promotions, products and printed materials for animals and get them for products or items in the shopping cart that must always be done by a local team of technicians. Including customer information. It is important to the customers by checking their address information. Or if you mean on the help page, we'll get it.

Table 11: Problematic examples.

Arch.	PPL	DIST/ENT \uparrow	S-BL \downarrow	Rep-4 \downarrow	BL \uparrow	R-L \uparrow	Score \uparrow	Len
DDIM vs DDPM								
DDIM	47.36	0.17/6.602	0.52	1.55%	0.77	7.9	0.55	168.08
DDPM	57.34	0.18/6.663	0.44	1.48%	0.35	5.7	0.53	162.81
Different Inference Steps								
5 steps	53.215	0.17/6.547	0.54	1.81%	0.67	7.2	0.55	134.2
10 steps	47.807	0.17/6.580	0.5	1.60%	0.69	7.4	0.55	138.78
20 steps	47.559	0.17/6.581	0.52	1.57%	0.71	7.7	0.55	146.38
30 steps	47.36	0.17/6.602	0.52	1.55%	0.77	7.9	0.55	168.08
50 steps	47.096	0.17/6.605	0.54	1.56%	0.83	7.9	0.55	162.08
Scheduler								
Cosine	47.36	0.17/6.602	0.52	1.55%	0.77	7.9	0.55	168.08
Beta Linear	49.78	0.17/6.577	0.57	1.46%	0.64	6.8	0.53	158.83

Table 12: Ablations on hotel review completion.

