

SongRewriter: A Chinese Song Rewriting System with Controllable Content and Rhyme Scheme

Yusen Sun^{1,2,*}, Liangyou Li², Qun Liu², Dit-Yan Yeung¹

¹The Hong Kong University of Science and Technology, Hong Kong SAR, China

²Huawei Noah's Ark Lab

ysunbc@connect.ust.hk, {liliangyou, liuqun}@huawei.com, dyyeung@cse.ust.hk

Abstract

Although lyrics generation has achieved significant progress in recent years, it has limited practical applications because the generated lyrics cannot be performed without composing compatible melodies. In this work, we bridge this practical gap by proposing a song rewriting system which rewrites the lyrics of an existing song such that the generated lyrics are compatible with the rhythm of the existing melody and thus singable. In particular, we propose *SongRewriter*,¹ a controllable Chinese lyrics generation and editing system which assists users without prior knowledge of melody composition. The system is trained by a randomized multi-level masking strategy which produces a unified model for generating entirely new lyrics or editing a few fragments. To improve the controllability of the generation process, we further incorporate a keyword prompt to control the lexical choices of the content and propose novel decoding constraints and a vowel modeling task to enable flexible end and internal rhyme schemes. While prior rhyming metrics are mainly for rap lyrics, we propose three novel rhyming evaluation metrics for song lyrics. Both automatic and human evaluations show that the proposed model performs better than the state-of-the-art models in both contents and rhyming quality.

1 Introduction

With the rapid development in natural language processing (Sutskever et al., 2014; Bahdanau et al., 2015; Vaswani et al., 2017; Devlin et al., 2019), lyrics generation has achieved significant advancement in recent years (Potash et al., 2015; Nikolov et al., 2020; Ma et al., 2021; Xue et al., 2021). Prior works mainly focus on two research topics:

*Work done during internship at Huawei Noah's Ark Lab.

¹Source code implemented in *MindSpore Lite* tool is available at <https://github.com/huawei-noah/noah-research/tree/master/NLP/SongRewriter>.

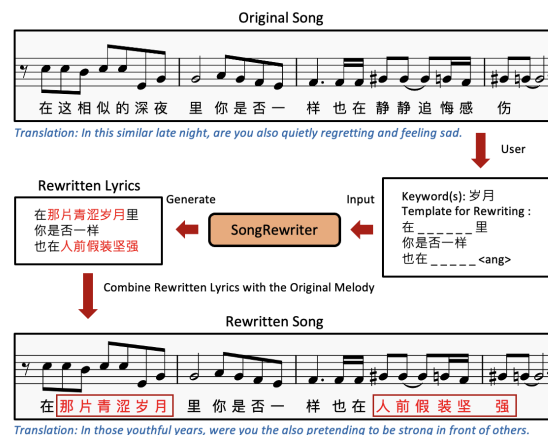


Figure 1: An overview of the proposed Chinese song rewriting system. Given an existing song, the users first mask the part(s) of the lyrics that they want to rewrite. Then, *SongRewriter* generates new lyrics corresponding to the masked fragments. Last, the rewritten lyrics are combined with the original melody to form a new song. During the generation, user can require the content to include specific keywords or control the rhyme scheme by setting the vowels of the output characters at specific positions.

singability of the generated lyrics (Watanabe et al., 2018; Lee et al., 2019) and controllability of the content and/or rhyme (Nikolov et al., 2020; Xue et al., 2021).

Current methods of generating singable lyrics is typically conditioned on a given melody. Accordingly, they treat the generation as a sequence-to-sequence translation task. However, there are two main challenges: 1) the parallel lyrics-melody dataset for training the model is limited, which mainly consists of 7,998 songs proposed by Yu et al. (2020); 2) melody notes and lyric syllables are loosely correlated and thus the alignments between them are hard to learn from the limited data. Therefore, prior works simplify the task by assuming a one-to-one mapping between the melody notes and

lyric syllables (Yu et al., 2020; Ma et al., 2021). However, such an assumption and restrictions may lead to a sub-optimal performance, as the mapping relationship is usually many-to-one in real life.

Another important research question on lyrics generation is how to control the generated content and rhyme. In prior works, the controllable content is usually achieved by conditioning the generation on given hint words or sentences (Shen et al., 2019; Zhang et al., 2022). However, they usually ignore the requirements of generating lyrics from a draft where the model needs to edit some sentences or words. In addition, prior works on rhyme control mainly focus on rhymes at the end of sentences (*end rhyme*) (Potash et al., 2015; Nikolov et al., 2020; Xue et al., 2021; Liu et al., 2022). To the best of our knowledge, no work has been done on both *internal rhyme* and *end rhyme* schemes.

In this work, we develop a user-assist AI system *SongRewriter* which can generate singable lyrics by rewriting parts of or the whole lyrics of a given song (i.e., *partial rewriting* and *full rewriting*, respectively) with controllable contents and rhyme schemes, as illustrated in Figure 1. To address the difficulty of learning the correlation between the melody and lyrics from a limited amount of parallel dataset, we propose to generate lyrics which have the same number of syllables as the original lyrics of the song. Therefore, the generated lyrics can well align with the melody. In addition, this method can directly learn the generation from text data so as to bypass the demand of parallel datasets.

Specifically, we adopt a transformer-based sequence-to-sequence auto-regressive model (Vaswani et al., 2017) as our model backbone. The model is trained by masking a few random fragments from the encoder’s inputs and predicting the masked fragments by the decoder. The generation process can be controlled in terms of three aspects. 1) To enable rewriting arbitrary parts of the lyrics, we train the model by randomly performing one of the three masking (i.e., token-level, sentence-level and song-level) strategies, corresponding to different levels of rewriting tasks. 2) To control the lexical choice of contents, we allow the generation to condition on given keyword prompts. This is achieved by training on extracted keywords from masked positions as additional encoder inputs. 3) To enable lyrics generation with arbitrary pre-defined rhyme schemes with vowel specified, we introduce additional vowel inputs and apply a

vowel mask strategy during training. This equips the model with the ability of generating tokens with required vowels at arbitrary positions. Since end rhyme is the most frequently used rhyme type, we specifically incorporate reverse language modelling and propose decoding constraints during inference to improve the vowel consistency and rhyming word diversity.

We evaluate our model on both generation controllability and quality in terms of keyword recall, vowel accuracy, lexical diversity, coherence, perplexity and rhyme quality. Since prior evaluation metrics on rhyme quality are mainly for rap lyrics and ignore the problem of identical rhyming words, we propose three new rhyme metrics to measure the local rhyme between adjacent sentences, global rhyme and diversity of rhyming words, respectively. Experimental results show that our model outperforms baseline models and state-of-the-art models on both full and partial rewriting tasks.

Our contributions are summarized as follows:

- We propose *SongRewriter* which generates melody-aligned lyrics by rewriting the lyrics of songs. It bypasses the difficulties of modelling the melody-lyrics correlation from limited parallel datasets.
- We propose a multi-level randomized masking scheme for training *SongRewriter*, which allows the model to rewrite arbitrary parts of the inputs according to the bidirectional context.
- We introduce a partial vowel masking strategy into training to enable lyrics generation on any rhyme schemes, and we design a novel decoding strategy to improve the end rhyme consistency and rhyming word diversity.
- We propose novel metrics for rhyme evaluation. We collect data from the internet for training and testing. Experiments show the effectiveness of our proposed model.

2 Related Work

Lyrics generation can be mainly divided into two categories, rap lyrics generation (Potash et al., 2015; Malmi et al., 2016; Nikolov et al., 2020; Xue et al., 2021) and song lyrics generation (Ramakrishnan A et al., 2009; Watanabe et al., 2018; Fan et al., 2019; Lee et al., 2019; Sheng et al., 2021; Ma et al., 2021; Liu et al., 2022). While rap lyrics generation

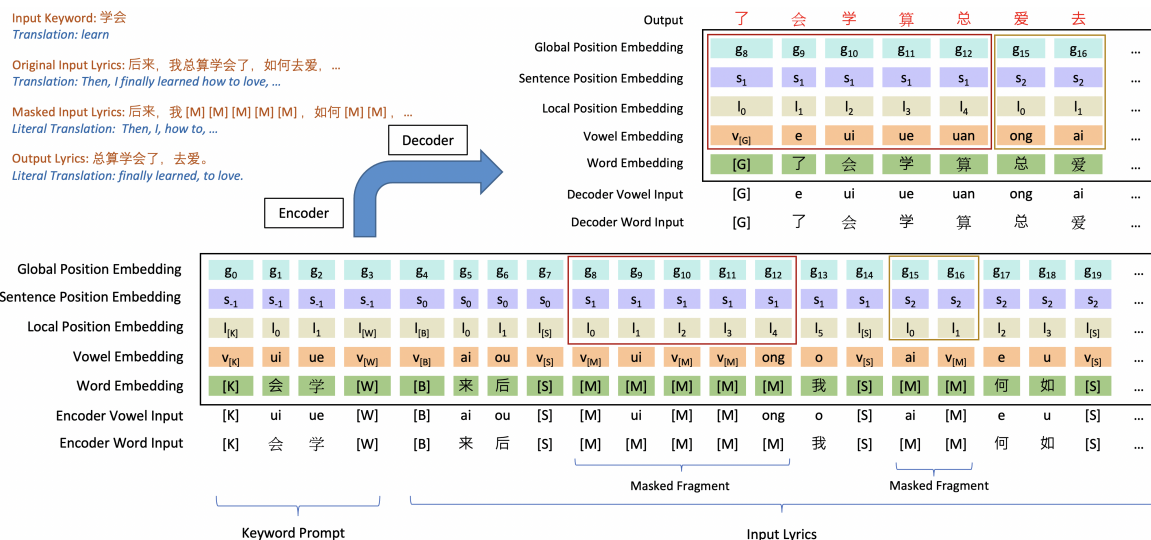


Figure 2: The architecture of the proposed *SongRewriter* model. The inputs of the encoder consist of a keyword prompt and partially masked lyrics. The keywords in the prompt are extracted from the masked fragments during training. The decoder uses [G] as a start token and generates the masked tokens autoregressively. The text order is reversed for end rhyme modeling. The original lyrics is from *Later* by *Rene Liu*.

mainly focuses on improving the rhyming performance (Nikolov et al., 2020; Xue et al., 2021), song lyrics generation concentrates on the alignment between melody and lyrics (Watanabe et al., 2018; Lee et al., 2019; Sheng et al., 2021).

Rhyming is an essential element for lyrics and poetry. To model rhyme, current approaches can be mainly divided into three types. The first type is to encourage the model to learn the rhyme structure implicitly during the training by adding additional rhyme signals. Potash et al. (2015) append $\langle endLine \rangle$ token to each sentence. Li et al. (2020) employ an additional format embedding as inputs to emphasize the rhyming tokens. Lau et al. (2018) incorporate an extra model to encode the ending tokens of the sentences. While Zhang et al. (2020) and Xue et al. (2021) both generate the rhyming word before the rest of the sentence, Zhang et al. (2020) move the last word of the sentence to the front, and Xue et al. (2021) generate sentence from right to left by reversing the word order. The second type is to apply explicit rhyme constraints during the training to force the model to generate rhyming sentences. Jhamtani et al. (2019) apply a discriminator on the sentence-ending words to learn the rhyme pattern adversarially. The last type is through post-editing, where the model first generates the lyrics, then another model edits the ending words to fulfil the rhyming constraints (Nikolov et al., 2020).

Most prior works focus on the rhyme at the end of the sentences (Xue et al., 2021; Liu et al., 2022). In this work, we extend the rhyme control to arbitrary rhyme schemes for both internal rhyme and end rhyme. To the best of our knowledge, this is the first work enabling arbitrary rhyme schemes.

3 Method

The proposed model is a transformer-based autoregressive sequence-to-sequence model (Vaswani et al., 2017). As shown in Figure 2, given an input of masked lyrics and keyword prompt to the encoder, the decoder generates output tokens corresponding to the masked tokens of the input, which contain the keywords in the prompt and satisfy the vowel constraints. Such generation also improves decoding efficiency and forces the model to rely more on the source input.

In a basic setting, given the lyrics of a song, the tokens of the input lyrics are of the following form,

$$[B], x_{00}, \dots, [M], \dots, [S], \dots, x_{ij}, \dots, [S], [E]$$

where x_{ij} denotes the j th tokens in the i th sentence of the lyrics, $[S]$ is the inter-sentence delimiter, $[B]$ is placed at the beginning of the lyrics with $[E]$ at the end, and the tokens to be rewritten are replaced by $[M]$ which will be predicted by the decoder.

To enable the controllability of the generation process, we further incorporate the keyword

prompt into the model to control the lexical choices, propose the multi-level masking strategy to enable rewriting arbitrary parts of the input in a single model and the rhyme control strategy to inject pre-defined rhyme schemes and improve rhyming word diversity.

3.1 Keyword Prompt

Using keyword prompts to control the text generation has been explored in other tasks, such as poetry generation (Zhipeng et al., 2019). In the task of lyrics generation, prior works mainly focus on theme control (Shen et al., 2019). In this work we introduce the technique of keyword prompts into lyrics generation to force the model generating lyrics containing these keywords.

Specifically, a keyword prompt is prepended to the input lyrics. The keyword prompt is a concatenation of a set of keywords in the following format,

$$[K], k_{00}, \dots, [W], k_{10}, \dots, [W], \dots, k_{ij}, \dots, [W]$$

where k_{ij} is the j th token in the i th keyword, $[W]$ is the inter-keyword delimiter, and $[K]$ is the start token of the prompt.

During training, we first use *jieba*² to segment the masked fragments into words and obtain their Parts-of-Speech tags. Then, we use the nouns and verbs to form a keyword database. Last, we sample a random number of keywords (ranging from 0 to 5) from the keyword database to form a keyword prompt. During inference, the keyword prompt is optional and can be provided by users.

3.2 Randomized Multi-Level Masking Scheme

As song rewriting requires the number of syllables between the original lyrics and the generated lyrics to be identical, we adopt the framework of MASS (Song et al., 2019), a sequence-to-sequence model pre-training method, which uses the decoder to predict the masked tokens in the encoder. However, MASS trains the model on inputs of single-sentence examples by masking a fragment of continuous tokens (around 50% of the input). Therefore, it is not an optimal strategy for full lyrics generation and editing arbitrary parts of the input lyrics.

Accordingly, we propose a novel masking strategy which masks the input from three levels (token-level, sentence-level and song-level) to simulate the

partial rewriting and full rewriting tasks. During training, for each input lyric, we randomly apply one of the following masking strategies:

- **Token-Level Masking:** To simulate the task of rewriting phrases in a sentence, for each sentence in the input lyric, we mask a few fragments of the sentence with a random ratio and train the model to reconstruct the masked portions of the sentences.
- **Sentence-Level Masking:** To enable sentence rewriting, we mask a random ratio of sentences (entire sentences) from the input and train the model to reconstruct the masked sentences.
- **Song-Level Masking:** We masks all the input tokens to simulate the full rewriting task.

We denote the above three masking schemes as {TOKEN, SENT, ALL}, respectively. During training, for SENT, we sample a masking ratio from a uniform distribution $U(0, 1)$ and randomly select the corresponding ratio of sentences. For TOKEN, we sample a masking ratio from a uniform distribution $U(0, 1)$ for each sentence and then randomly select the ratio of tokens to mask.

3.3 Rhyme Modeling and Control

Our method for rhyme modelling and control is divided into two parts, rhyming modelling and control for the final syllables of the lines (*end rhyme*), and rhyme control for an arbitrary rhyme scheme which defines the required vowels at specific positions in the generated lyrics (*internal rhyme*).

End Rhyme Modeling and Constraint

End rhyme is the most frequently used rhyme type for lyrics and poems. It occurs when the last words of the sentences rhyme. Inspired by the rhyme modelling for rap lyrics (Xue et al., 2021), we adopt reverse language modelling with two additional position embeddings, sentence position embedding and local position embedding, to facilitate the modelling of rhyme features in the lyrics. Specifically, we reverse the order of the characters in each sentence for both inputs and target outputs (while keeping the sentence order unchanged). Therefore, the reverse sentence starts with the potential rhyming character, i.e., the end character in the original sentence. Accordingly, the model can easily learn to identify the rhyming characters from the inputs with the local position l_0 .

²<https://github.com/fxsjy/jieba>

However, since rhyming with identical words is considered inferior, we incorporate two control factors at inference time to encourage rhyme consistency and rhyming word diversity. Given the end character set $e_{<t} = \{e_0, \dots, e_{t-1}\}$ and their corresponding vowel set $v_{<t} = \{v_0, \dots, v_{t-1}\}$ from the previous t sentences, we define an adjusted probability of the end character of the $(t + 1)$ th sentence being x_i as,

$$\bar{p}_t^i = \frac{p_t^i \cdot \Lambda(v_{x_i}, v_{<t}) \cdot \Gamma(x_i, e_{<t})}{\sum_j^N p_t^j \cdot \Lambda(v_{x_j}, v_{<t}) \cdot \Gamma(x_j, e_{<t})} \quad (1)$$

where p_t^i is the predicted probability of the token x_i in the vocabulary, and v_{x_i} is the vowel of the token x_i . The two factors:

- $\Lambda(\cdot)$ returns λ if the vowel of the token x_i appears in the end vowel set $v_{<t}$ of previous sentences; otherwise, returns 1.
- $\Gamma(\cdot)$ returns γ if x_i appears in the end character set $e_{<t}$ of previous sentences, otherwise 1.

λ and γ are the hyper-parameters to control the rhyming effect. While a larger λ increases the probability of the model sampling a word with the same vowel in the previous end vowel set (improving rhyme consistency), a smaller γ reduces the chance of a generated end word being chosen again (increasing rhyming word diversity).

Internal Rhyme by Vowel Modeling

While rhyme generally refers to end rhyme, where the last words of the lines rhyme with each other, internal rhyme has also been widely used, which usually includes multiple rhyming words either within the same line (usually one in the middle and the other one at the end) or in the middle of multiple lines. As shown in Figure 3, the highlighted characters (with the corresponding *pinyin* in parentheses) share the same vowel and rhyme. While end rhyme has been widely investigated for both poetry generation (Lau et al., 2018; Li et al., 2020) and rap lyrics generation (Xue et al., 2021; Nikolov et al., 2020), internal rhyme is still yet to be explored.

To model internal rhyme as well as other rhyme schemes with specified vowels, we propose a *restricted vowel loss* to enable direct control of vowels at arbitrary positions. Specifically, as shown in Figure 2, during the training stage, for the masked

色白花青的锦鲤(ii)跃然于碗底(di),
On the bottom of the white bowl, blue and white carps leap to and fro.
临摹宋体(ti)落款时却惦记着你(ni),
Copying the Song inscription, to the end I am yet to sign, before you step into my mind.

Figure 3: An example of lyrics with internal rhyme. The lyrics are from *Blue and White Porcelain*, a Chinese song of Jay Chou, with the English translation underneath. The rhyming characters are highlighted in red with their *pinyin*.

fragments in the input, we only replace 80% of the vowel inputs with the $[M]$ token. For the remaining 20% masked input tokens with ground truth vowel inputs, we introduce an additional vowel prediction task. For the t th output token with a predicted distribution $p_t \in \mathcal{R}^{N \times 1}$, if the ground truth vowel v_t is provided in the input, the additional training objective is

$$\mathcal{L}_t^v = -\log p(v_t | y_{<t}, \text{enc}(x_{1:T})) \quad (2)$$

where the probability of the vowel v_t at the time step t is calculated by,

$$p(v_t | \cdot) = \sum_i^N p_t^i \cdot \mathbb{1}(V(x_i) = v_t) \quad (3)$$

where p_t^i is the predicted probability for each token i in the vocabulary at the decoding time step t , and $V(x_i)$ is a mapping function which returns the vowel of the token x_i . The function $\mathbb{1}(V(x_i) = v_t)$ returns 1 if the vowel of the token x_i is identical to the ground truth vowel v_t , otherwise returns 0. Thus, $p(v_t | \cdot)$ is basically summing up the predicted probabilities of all the tokens with the same vowel v_t .

During the inference stage, the internal rhyme scheme can be created for the generated outputs by providing rhyming vowel inputs to specific positions in masked fragments of the inputs.³

4 Experiment

4.1 Datasets and Baselines

Our model is trained on three different datasets. We first crawl a large-scale text corpus from *Baidu Encyclopedia*, which is the largest Chinese online Encyclopedia. We use this dataset to pretrain our model. Then, we crawl a lyrics dataset from two Chinese lyrics websites.⁴ Since the amount of the

³An example is shown in Figure 8 in Appendix J.

⁴<https://mojim.com> and <https://www.mulanci.org>

lyrics data is limited, we further crawl proeses as a supplementary dataset from a Chinese prose website.⁵ The pretrained model is first fine-tuned on the prose dataset and then on the the lyrics dataset to produce our final system.⁶

Since Chinese is a monosyllabic language, each character consists of one syllable. To control the number of syllables in the generated output, We use the *BasicTokenizer* from the *transformers* library (Wolf et al., 2020) for tokenization, which splits text into characters for Chinese and into words for other languages (mainly English). We keep the words and characters with frequency larger than 3,000 to build a vocabulary of size 6,572. For the vowels, we employ *python-pinyin*⁷ to extract the vowels of the Chinese characters. There are in total 21 distinct vowels.

We evaluate the proposed model, *SongRewriter*, for both full and partial rewriting tasks. For the full rewriting task, we compare *SongRewriter* with a Chinese *GPT2* (Radford et al., 2019) and *SongNet* (Li et al., 2020). For the partial rewriting task, we compare the proposed model with *ILM* (Donahue et al., 2020). We incorporate the keyword prompt function to the *ILM* model resulting in *ILM-Keyword* which is used as a baseline model for keyword-conditioned rewriting tasks. All the baseline models are fine-tuned on our lyrics dataset for a fair comparison.⁸

4.2 Evaluation Metrics

We evaluate the performance in terms of two aspects, generation controllability and generation quality. The controllability metrics include *Keyword Recall* and *Vowel Accuracy*. The quality metrics include *Diversity*,⁹ *Coherence*,¹⁰ *Perplexity-Test (PPL-Test)* and *Perplexity-Gen (PPL-Gen)*.¹¹ Following prior works, we assume the optimal approach should generate lyrics with quality closest to the human-written lyrics (Holtzman et al., 2020). Therefore, we report the absolute difference scores on some metrics: $\Delta Diversity$, $\Delta Coherence$ and $\Delta PPL-Gen$.

To measure rhyme quality, we propose three novel metrics:

⁵<https://www.sanwenwang.com/sanwen/>

⁶Details of the datasets are in Table 5 in Appendix A.

⁷<https://github.com/mozillazg/python-pinyin>

⁸The details of the baselines are in the Appendix B.

⁹*Diversity* is evaluated by *distinct* (Li et al., 2016).

¹⁰We measure *Coherence* based on sentence similarity.

¹¹Details on these metrics can be found in Appendix D.

- **Local Rhyme (Rhyme-L):** While a lyric may contain multiple rhyming vowels, sentences sharing the same rhyme are usually grouped together. Therefore, we propose *local-rhyme-n* to evaluate this localised characteristics. Specifically, We define a sentence being *locally n-rhymed* if, among the n sentences before and after the current sentence, there are at least one sentence sharing the same rhyming vowel with the current sentence. Thus, *local-rhyme-n* is defined as the number of locally n -rhymed sentences divided by total number of sentences. We report *Rhyme-L*, which is the average of the *local-rhyme-n* with $n \in [1, 4]$.
- **Global Rhyme (Rhyme-G):** Apart from evaluating rhyming effect from a local perspective, we also measure the rhyming effect of a lyrics as a whole. Since the more sentences sharing the same rhyming vowels, the better the rhyming effect. We evaluate the global rhyming performance by calculating the portion of duplicated rhyming vowels: $1 - \frac{\text{number of unique rhyming vowels}}{\text{total number of sentences}}$.
- **Diversity of Rhyming Words (Dist-RW):** Since rhyming with identical words is considered inferior, we evaluate the diversity of the rhyming words by calculating the ratio of the number of unique end words to the total number of end words.

4.3 Results and Discussion

Full Song Rewriting

We evaluate the performance of *SongRewriter* on the task of full song rewriting by masking all the input tokens. This task is similar to lyrics generation with a fixed format, i.e., the number of sentences and lengths of each sentence are pre-defined.

As shown in Table 1, *SongRewriter* significantly outperforms other models in terms of both rhyme quality and content quality. For the content, the generated outputs of *SongRewriter* are closer to the human-written lyrics in the aspects of lexical diversity, coherence and fluency (*PPL-Gen*), and *PPL-Test* further verifies the language modeling ability of *SongRewriter*.

By inspecting the generated outputs,¹² we observe that: 1) The generated lyrics are fluent and coherent in general. 2) Most lines share the same vowel and rhyme with their neighbouring

¹²Figure 5 in Appendix J shows an example.

Model	Δ Diversity \downarrow	Δ Coherence \downarrow	Δ PPL-Gen \downarrow	PPL-Test \downarrow	Dist-RW \uparrow	Rhyme-L \uparrow	Rhyme-G \uparrow
GPT2	0.135	0.076	3.695	3.539	0.437	0.546	0.723
SongNet	0.110	0.046	2.569	3.429	0.465	0.605	0.766
SongRewriter	0.023	0.040	1.287	3.157	0.552	0.737	0.812

Table 1: Quality evaluation on the generated outputs in the full song rewriting task. The best scores are in bold.

Mask	Model	Δ Diversity \downarrow	Δ Coherence \downarrow	Δ PPL-Gen \downarrow	Dist-RW \uparrow	Rhyme-L \uparrow	Rhyme-G \uparrow
TOKEN	ILM	0.265	0.052	17.88	0.547	0.658	0.762
	SongRewriter	0.039	0.006	2.37	0.572	0.722	0.800
SENT	ILM	0.057	0.024	0.70	0.500	0.658	0.781
	SongRewriter	0.043	0.002	1.66	0.578	0.742	0.814

Table 2: Quality evaluation on the generated outputs on the task of partial song rewriting under masking schemes, {SENT, TOKEN}. We report the averaged scores over three masking ratios, {0.25, 0.5, 0.75}. The best scores are in bold.

lines with diverse rhyming words. 3) Similar to human-written lyrics, automatically generated lyrics contain duplicated blocks, which indicates that *SongRewriter* can learn the structural information of the lyrics.

Partial Song Rewriting

We test the performance of *SongRewriter* on the partial song rewriting task by masking a portion of the input lyrics. We compare the proposed model with *ILM* under two masking schemes, {SENT, TOKEN}. We average the scores under three masking ratios, {0.25, 0.5, 0.75}.¹³

As shown in Table 2, the proposed model performs better than *ILM* in general at both token-level and sentence-level masking. Specifically, for *ILM*, it is observed that there is a large discrepancy on content quality (*Diversity*, *Coherence* and *PPL-Gen*) between the TOKEN and SENT masking schemes. On the contrary, *SongRewriter* achieves consistent performance, indicating that the proposed method is more suitable for both tasks. Besides, we also find that the performance of *ILM* decreases significantly as the masking ratio increases, while *SongRewriter* consistently achieves high performance across various masking ratios. This suggests that *SongRewriter* is able to tackle arbitrary portion of the rewriting consistently. By inspecting examples,¹⁴ we observe that, during the partial song rewriting, *SongRewriter* considers not only the bidirectional context but also the rhyming effects with the input sentences.

Mask	Model	Keyword Recall \uparrow
TOKEN	ILM-Keyword	0.545
	SongRewriter	0.890
SENT	ILM-Keyword	0.451
	SongRewriter	0.877
ALL	ILM-Keyword	0.851
	SongRewriter	0.952

Table 3: Evaluation results on the task of keyword-conditioned partial song rewriting under masking schemes, {SENT, TOKEN, ALL}. We report the averaged scores over three masking ratios, {0.25, 0.5, 0.75}. The best scores are in bold.

Keyword Control

We test the performance of controlled lyrics rewriting under masking schemes {SENT, TOKEN}. We report the averaged scores over three masking ratios, {0.25, 0.5, 0.75}.

To evaluate the ability of generating content with keywords, we first build a keyword database by using *jieba*¹⁵ to extract keywords from the training set. We evaluate the controllability by sampling keywords from the database.¹⁶ As shown in Table 3, *SongRewriter* performs significantly better than the baseline model, *ILM-Keyword*, by a large margin.

Rhyme Scheme Control

To evaluate the rhyme scheme controllability, we evaluate the ability of the model to generate tokens with pre-defined vowels at arbitrary positions. We randomly mask 80% of the vowel inputs from the

¹³Detailed results are in Table 7 in Appendix F.

¹⁴Figures 8 and 9 in Appendix J.

¹⁵<https://github.com/fxsjy/jieba>

¹⁶We define the sampling probability of keyword k_i to be the number of occurrence of k_i divided by the number of occurrence of all keywords.



Figure 4: Ablation results on the generated outputs on the task of full song rewriting. To facilitate comparison, all metrics are normalized to 0-1 range by the respective maximum values. The exact scores are presented in Table 8 in Appendix H.

masked fragments. For the remaining 20% masked tokens with vowel inputs, we calculate the ratio of output tokens with the same vowel as the input (*vowel accuracy*). We found that the proposed model is able to consistently generate tokens of pre-defined vowels around 98% of the time under various masking schemes and masking ratios,¹⁷ implying that the model can generate user-defined rhyme schemes by providing rhyming vowel to the target positions in the input most of the times.

Ablation Study

As shown in the Figure 4, by removing the inference constraints, the performance on *Rhyme-L* and *Rhyme-R* increases while the others decrease. By inspecting the output samples, it is found that the model without inference constraints is more likely to generate repetitive sentences and rhyming words, thus leading to an increase in *Coherence* (more similar content), *Rhyme-L* and *Rhyme-R*, but a decrease in content diversity and rhyming word diversity.

To verify the effectiveness of the multi-level masking scheme, we train a model with the masking scheme proposed in MASS (Song et al., 2019), which is to mask 50% of the tokens consecutively from the inputs. As shown in Figure 4, the language fluency and the rhyming performance drop significantly in terms of $\Delta PPL-Gen$, *Rhyme-L* and *Rhyme-G*. Although the diversity of the rhyming words increases, those ending words do not share the same vowel, and thus, not rhyme with each other. The performance drop is expected as by only masking 50% of tokens from the inputs, there is a

¹⁷Detailed results are in Table 9 in Appendix G.

Mask	Model	Fluency	Coherence	Rhyme
TOKEN	ILM	2.63	2.89	2.58
	SongRewriter	3.58	3.34	3.75
SENT	ILM	3.36	2.96	2.91
	SongRewriter	3.71	3.17	3.82
ALL	GPT2	3.31	3.18	2.56
	SongNet	3.59	3.25	3.11
	SongRewriter	3.76	3.38	3.93

Table 4: Human evaluation results on the generated outputs. For the partial song rewriting, the masking ratio is set to 0.5.

task discrepancy between the training task (rewriting half of the tokens) and inference task (rewriting content with ratios ranging from 0 to 1).

Regarding the vowel modeling, by removing the vowel loss from the training objective, the vowel accuracy for the full song rewriting drops from 98.5% to 92.5%, showing that incorporating vowel loss can help the model generate tokens with correct vowels at the specific positions.

Human Evaluation

Apart from automatic metrics, we also conduct human evaluation following the previous works (Lee et al., 2019; Xue et al., 2021). We sample 200 examples from the test set as inputs and generate 200 outputs from each model. Then, we recruit 3 annotators with musical knowledge background to score the generated lyrics from 1 (Poor) to 5 (Perfect) on three criteria: language fluency, content coherence and rhyme quality.

As shown in Table 4, the human evaluation results show that *SongRewriter* outperforms other models on all three tasks (full song rewriting, sentence rewriting and partial sentence rewriting). In particular, *SongRewriter* performs significantly better in terms of rhyming.

5 Conclusion

In this work, we propose to overcome the difficulties of modelling the melody-lyrics correlation from limited parallel datasets by directly rewriting the lyrics of songs. We propose a unified model for full and partial song rewriting by training with a multi-level randomized masking scheme. The proposed model allows rewriting arbitrary parts of the inputs according to the bidirectional context. Besides, we introduce a partial vowel masking strategy into training to enable lyrics generation on any rhyme schemes. A novel decoding strategy is designed to improve the end rhyme consistency and rhyming word diversity. Novel metrics are pro-

posed for rhyme evaluation. Both automatic and human evaluation shows our proposed model outperforms baseline and state-of-the-art models.

6 Limitation

Since each Chinese character contains 1 syllable, our proposed model can control the number of syllables in the generation by the number of generated tokens. However, this method does not apply to languages with multisyllabic words (such as English). To rewrite lyrics with multisyllabic words while maintaining the same number of syllables, a special technique such as syllable-level subword tokenization may be needed. This line of work will be left to be investigated in the future.

7 Ethics Statement

Rewriting the lyrics of a song may cause potential copyright infringement. Besides, the copyrights of the lyrics in the dataset belong to the song writers. To protect the copyrights, our model and the released dataset will be protected by the license, Creative Commons Attribution-NonCommercial (CC-BY-NC), and prohibited from commercial use.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Chris Donahue, Mina Lee, and Percy Liang. 2020. [Enabling language models to fill in the blanks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2492–2501, Online. Association for Computational Linguistics.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. [Hierarchical neural story generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia. Association for Computational Linguistics.
- Haoshen Fan, Jie Wang, Bojin Zhuang, Shaojun Wang, and Jing Xiao. 2019. [A hierarchical attention based seq2seq model for chinese lyrics generation](#). *ArXiv*, abs/1906.06481.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text de-generation](#). In *International Conference on Learning Representations*.
- Harsh Jhamtani, Sanket Vaibhav Mehta, Jaime Carbonell, and Taylor Berg-Kirkpatrick. 2019. [Learning rhyming constraints using structured adversaries](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6025–6031, Hong Kong, China. Association for Computational Linguistics.
- Jey Han Lau, Trevor Cohn, Timothy Baldwin, Julian Brooke, and Adam Hammond. 2018. [Deep-speare: A joint neural model of poetic language, meter and rhyme](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1948–1958, Melbourne, Australia. Association for Computational Linguistics.
- Hsin-Pei Lee, Jhih-Sheng Fang, and Wei-Yun Ma. 2019. [iComposer: An automatic songwriting system for Chinese popular music](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 84–88, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A diversity-promoting objective function for neural conversation models](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Piji Li, Haisong Zhang, Xiaojiang Liu, and Shuming Shi. 2020. [Rigid formats controlled text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 742–751, Online. Association for Computational Linguistics.
- Nayu Liu, Wenjing Han, Guangcan Liu, Da Peng, Ran Zhang, Xiaorui Wang, and Huabin Ruan. 2022. [Chip-Song: A controllable lyric generation system for Chinese popular song](#). In *Proceedings of the First Workshop on Intelligent and Interactive Writing Assistants (In2Writing 2022)*, pages 85–95, Dublin, Ireland. Association for Computational Linguistics.

- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Xichu Ma, Ye Wang, Min-Yen Kan, and Wee Sun Lee. 2021. *AI-Lyricist: Generating Music and Vocabulary Constrained Lyrics*, page 1002–1011. Association for Computing Machinery, New York, NY, USA.
- Eric Malmi, Pyry Takala, Hannu Toivonen, Tapani Raiko, and Aristides Gionis. 2016. [Dopelearning: A computational approach to rap lyrics generation](#). In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, page 195–204, New York, NY, USA. Association for Computing Machinery.
- Nikola I. Nikolov, Eric Malmi, Curtis Northcutt, and Loreto Parisi. 2020. [Rapformer: Conditional rap lyrics generation with denoising autoencoders](#). In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 360–373, Dublin, Ireland. Association for Computational Linguistics.
- Peter Potash, Alexey Romanov, and Anna Rumshisky. 2015. [GhostWriter: Using an LSTM for automatic rap lyric generation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1919–1924, Lisbon, Portugal. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Ananth Ramakrishnan A, Sankar Kuppan, and Sobha Lalitha Devi. 2009. [Automatic generation of Tamil lyrics for melodies](#). In *Proceedings of the Workshop on Computational Approaches to Linguistic Creativity*, pages 40–46, Boulder, Colorado. Association for Computational Linguistics.
- Liang-Hsin Shen, Pei-Lun Tai, Chao-Chung Wu, and Shou-De Lin. 2019. [Controlling sequence-to-sequence models - a demonstration on neural-based acrostic generator](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP): System Demonstrations*, pages 43–48, Hong Kong, China. Association for Computational Linguistics.
- Zhonghao Sheng, Kaitao Song, Xu Tan, Yi Ren, Wei Ye, Shikun Zhang, and Tao Qin. 2021. [Songmass: Automatic song writing with pre-training and alignment constraint](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(15):13798–13805.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2019. [MASS: Masked sequence to sequence pre-training for language generation](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, page 5926–5936. PMLR.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: A simple way to prevent neural networks from overfitting](#). *Journal of Machine Learning Research*, 15(56):1929–1958.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'14, page 3104–3112, Cambridge, MA, USA. MIT Press.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Kento Watanabe, Yuichiroh Matsubayashi, Satoru Fukayama, Masataka Goto, Kentaro Inui, and Tomoyasu Nakano. 2018. [A melody-conditioned lyrics language model](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 163–172, New Orleans, Louisiana. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Lanqing Xue, Kaitao Song, Duocai Wu, Xu Tan, Nevin L. Zhang, Tao Qin, Wei-Qiang Zhang, and Tie-Yan Liu. 2021. [DeepRapper: Neural rap generation with rhyme and rhythm modeling](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 69–81, Online. Association for Computational Linguistics.
- Yi Yu, Florian Harscoët, Simon Canales, Gurunath Reddy M, Suhua Tang, and Junjun Jiang. 2020. [Lyrics-conditioned neural melody generation](#). In *MultiMedia Modeling: 26th International Conference, MMM 2020, Daejeon, South Korea, January 5–8, 2020, Proceedings, Part II*, page 709–714, Berlin, Heidelberg. Springer-Verlag.
- Le Zhang, Rongsheng Zhang, Xiaoxi Mao, and Yongzhu Chang. 2022. [QiuNiu: A Chinese lyrics generation system with passage-level input](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demon-*

strations, pages 76–82, Dublin, Ireland. Association for Computational Linguistics.

Rongsheng Zhang, Xiaoxi Mao, Le Li, Lin Jiang, Lin Chen, Zhiwei Hu, Yadong Xi, Changjie Fan, and Minlie Huang. 2020. [Youling: an AI-assisted lyrics creation system](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 85–91, Online. Association for Computational Linguistics.

Guo Zhipeng, Xiaoyuan Yi, Maosong Sun, Wenhao Li, Cheng Yang, Jiannan Liang, Huimin Chen, Yuhui Zhang, and Ruoyu Li. 2019. [Jiuge: A human-machine collaborative Chinese classical poetry generation system](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 25–30, Florence, Italy. Association for Computational Linguistics.

A Dataset Statistics

Model	#train	#dev	#test
Baidu Encyclopedia	15M	–	–
Prose	330k	2k	2k
Lyrics	154k	1k	1k

Table 5: Statistics on the number of documents, proses and lyrics. The dev and test sets are randomly sampled.

B Baselines

We evaluate the proposed model, *SongRewriter*, for both full and partial rewriting tasks. For the full rewriting task, we compare *SongRewriter* with *GPT2* (Radford et al., 2019) and *SongNet* (Li et al., 2020) which are fine-tuned on our lyrics datasets. For the partial rewriting task, we compare the proposed model with *ILM* (Donahue et al., 2020). The details of the baselines are as follows,

- **GPT2:** *GPT2* is an auto-regressive language model based on transformer decoder. Initializing with *GPT2-Chinese*,¹⁸ we fine-tune the model on the lyrics dataset. Note that the lyrics generated by *GPT2* is in free form, it does not follow any format constraints.
- **SongNet:** *SongNet* is a rigid format controlled text generation model which forces the generated output to follow the exact sentence lengths and sentence numbers of the input. We fine-tune their released pre-trained checkpoint¹⁹ on the lyrics dataset for full song rewriting.
- **ILM:** *ILM* is a GPT2-based model specialised on the text infilling task. It randomly replaces part of the text by "[blank]" tokens, and appends the masked segments (which are concatenated by "[answer]" tokens) to the end of the masked input text. For example, given a source text "*She ate leftover pasta for lunch*", an *ILM* example will be "*She ate [blank] for [blank] [sep] leftover pasta [answer] lunch [answer]*". We follow their training strategies and fine-tune *GPT2-Chinese* on the lyrics dataset for partial song rewriting.
- **ILM-Keyword:** We incorporate the keyword-controlled function to the *ILM* model by

¹⁸[uer/gpt2-chinese-cluecorpussmall](https://huggingface.co/openai/gpt2-chinese-cluecorpussmall)

¹⁹<https://github.com/lipiji/SongNet>

adding a keyword prompt to the *ILM* example, leading the input to be in the form, "[keyword] pasta [keyword] lunch [CLS] She ate [blank] for [blank] [sep] leftover pasta [answer] lunch [answer]". We use *ILM-Keyword* as a baseline model for keyword-conditioned tasks.

C Training and Inference Settings

The proposed model is a transformer encoder-decoder model (Vaswani et al., 2017). There are 12 layers for the encoder and decoder, respectively, with 12 heads for each layer. The hidden dimension is 768, and the dropout (Srivastava et al., 2014) is 0.1. We employ the AdamW (Loshchilov and Hutter, 2019) optimizer with a weight decay rate to be 10^{-4} . For the pre-training stage, we use 8,000 warm-up steps with the default learning rate schedule in Vaswani et al. (2017) and train for 10,000 iterations. On the fine-tuning stage, we use a fixed learning rate of 10^{-5} and train until models converge. We first fine-tune the model on the prose dataset. Afterwards, we fine-tune the model on the lyrics dataset. During inference, we apply *Top-K sampling* (Fan et al., 2018) with k to be 32. The rhyme factors, γ and λ , are set to 0.3 and 1.4, respectively.

D Definition of Evaluation Metrics

- **Keyword Recall:** To evaluate the content control ability by keyword prompt, we calculate the keyword recall rate, which is the percentage of the keywords appearing in the generated outputs.
- **Vowel Accuracy:** To measure the performance of the vowel control, we calculate the percentage of output tokens with correct vowels as the input.
- **Diversity:** We evaluate the diversity of the generated lyrics by *distinct-n* (Li et al., 2016), which is defined as the number of unique n-grams divided by total number of n-grams. We report the average of the *distinct-n* with n from 1 to 4:

$$\text{Diversity} = \frac{1}{4} \sum_{n=1}^4 \text{distinct-n}$$

- **Coherence:** To evaluate the semantic consistency among all the sentences of the generated lyrics, we measure the semantic textual

Model	Diversity	Local-STs	Global-STs	PPL-Gen	PPL-Test	Dist-RW	Rhyme-L	Rhyme-G
Test Set	0.584	0.276	0.269	6.192	NA	0.540	0.724	0.804
GPT2	0.449 _{0.005}	0.343 _{0.002}	0.353 _{0.001}	2.497 _{0.013}	3.539	0.437 _{0.005}	0.546 _{0.004}	0.723 _{0.005}
SongNet	0.474 _{0.002}	0.323 _{0.002}	0.313 _{0.001}	3.632 _{0.011}	3.429	0.465 _{0.002}	0.605 _{0.002}	0.766 _{0.001}
SongRewriter	0.561 _{0.002}	0.322 _{0.001}	0.302 _{0.001}	4.905 _{0.035}	3.157	0.552 _{0.002}	0.737 _{0.002}	0.812 _{0.001}

Table 6: Evaluation on the content quality of the generated lyrics in full song rewriting task. For each model, we apply *Top-K sampling* (Fan et al., 2018) for five times and report the mean and standard deviation (subscript).

Model	Mask	Ratio	Diversity	Local-STs	Global-STs	PPL-Gen	Dist-RW	Rhyme-L	Rhyme-G
Test Set	NA	NA	0.584	0.276	0.269	6.192	0.540	0.724	0.804
ILM	TOKEN	0.25	0.657	0.271	0.253	15.99	0.602	0.704	0.783
		0.5	0.668	0.292	0.270	22.72	0.594	0.654	0.762
		0.75	0.692	0.323	0.293	33.49	0.446	0.616	0.740
	SENT	0.25	0.588	0.281	0.269	7.03	0.548	0.686	0.785
		0.5	0.536	0.295	0.281	6.71	0.506	0.656	0.782
		0.75	0.467	0.333	0.319	5.44	0.446	0.633	0.777
SongRewriter	TOKEN	0.25	0.632	0.266	0.253	8.30	0.590	0.737	0.802
		0.5	0.612	0.275	0.265	8.22	0.564	0.710	0.796
		0.75	0.623	0.276	0.264	9.16	0.563	0.719	0.801
	SENT	0.25	0.624	0.277	0.264	7.38	0.566	0.718	0.800
		0.5	0.614	0.278	0.265	7.57	0.559	0.743	0.815
		0.75	0.644	0.280	0.265	8.60	0.608	0.766	0.826

Table 7: Evaluation results on the content quality of the generated outputs on the task of partial song rewriting under masking schemes, {SENT, TOKEN}, and masking ratios, {0.25, 0.5, 0.75}.

similarity (STS) between all sentence pairs. Specifically, we employ pre-trained *SimCSE* (Gao et al., 2021) to extract sentence embeddings and calculate the cosine similarity.²⁰

Global-STs =

$$\frac{2}{N(N-1)} \sum_{i=1}^{N-1} \sum_{j=i+1}^N \text{SimCSE}(s_i, s_j)$$

Besides, we also calculate the STS between adjacent sentences for the consistency from a local aspect:

$$\text{Local-STs} = \frac{1}{N-1} \sum_{i=1}^{N-1} \text{SimCSE}(s_i, s_{i+1})$$

We define the coherence to be the average:

$$\text{Coherence} = \frac{1}{2}(\text{Global-STs} + \text{Local-STs})$$

- **PPL-Test:** We evaluate the quality of the model by calculating the model perplexity on the test set.
- **PPL-Gen:** We evaluate the quality of the generated lyrics with the perplexity from a language model, which is a pre-trained Chinese

²⁰We use `sentence-Transformers` library with pre-trained model `cyclone/simcse-chinese-robetta-wwm-ext`.

GPT2 model²¹ (Radford et al., 2019) fine-tuned on the lyrics dataset.

E Full Song Rewriting

Table 6 shows evaluation results on the content quality of the generated lyrics in full song rewriting task. For each model, we apply *Top-K sampling* (Fan et al., 2018) for five times and report the mean and standard deviation (subscript). Figures 5 and 6 show two examples of the generated lyrics.

F Partial Song Rewriting

Table 7 shows evaluation results on the generated outputs on the task of partial song rewriting under masking schemes, {SENT, TOKEN}, and masking ratios, {0.25, 0.5, 0.75}. Figure 8 shows an example of token-level rewriting. Figure 9 shows an example of sentence-level rewriting.

G Rhyme Scheme Control

Table 9 shows the evaluation results on the task of vowel-conditioned partial song rewriting under the masking schemes, {SENT, TOKEN}, and masking ratios, {0.25, 0.5, 0.75}.

Figure 7 shows an example of applying rhyme scheme control. We first extract the syllable template and rhyme scheme from the original lyrics.

²¹We use the pre-trained model `uer/gpt2-chinese-cluecorpussmall`.

Model	Δ Diversity \downarrow	Δ Coherence \downarrow	Δ PPL-Gen \downarrow	Dist-RW \uparrow	Rhyme-L \uparrow	Rhyme-G \uparrow
SongRewriter	0.023	0.040	1.287	0.552	0.737	0.812
- inference Constraints	0.151	0.105	2.452	0.325	0.767	0.839
- Fixed Masking Ratio (0.5)	0.047	0.013	2.964	0.594	0.568	0.753

Table 8: Ablation results on the generated lyrics in full song rewriting task.

As shown in the middle column of Figure 7, the original lyrics control both *end rhyme* and *internal rhyme*. The end rhyme scheme is AAAA, where all four lines of a verse share the same ending vowel and rhyme with each other. As shown in Figure 7, by inputting the desired vowel tokens in the target positions, the generated lyrics have a consistent rhyme scheme with the original lyrics.

Mask	Ratio	Vowel Accuracy
TOKEN	0.25	0.984
	0.5	0.982
	0.75	0.985
SENT	0.25	0.984
	0.5	0.984
	0.75	0.984
ALL	1.0	0.985

Table 9: Evaluation results of *SongRewriter* on the task of vowel-conditioned partial song rewriting under masking schemes, {SENT, TOKEN}, and masking ratios, {0.25, 0.5, 0.75}.

H Ablation Study

Table 8 shows ablation results on the generated lyrics in full song rewriting task.

Model	Dist-RW	Rhyme-L	Rhyme-G
SongRewriter	0.552	0.737	0.812
- Seq TK & Seq LP	0.500	0.510	0.712
- Seq TK & Rev LP	0.531	0.577	0.745
- Rev TK & Rev LP	0.511	0.707	0.806

Table 10: Evaluation results on the rhyming performance of the generated lyrics on the full song rewriting task under different language order and local position embedding order.

I Sequence Order and Local Position Order

To improve *end rhyme* modeling, we incorporate *reverse language modeling* with *local position embeddings*. We assume that by reversing the language order, the model is easier to locate the rhyming words (which corresponds to " $\langle l_0 \rangle$ ", the first token of the reverse sentence according to the local position embedding) and generate rhyming

sentences by generating the end words before the rest of the sentences. In this section, we verify the effectiveness of "labeling rhyming words by " $\langle l_0 \rangle$ " and "generating rhyming word of the sentence first" by comparing four model variants: *reverse token order with sequential local position* (the proposed method), *reverse token order with reverse local position* (*Rev TK & Rev LP*), *sequential token order with sequential local position* (*Seq TK & Seq LP*) and *sequential token order with reverse local position* (*Seq TK & Rev LP*).

As shown in Table 10, the rhyming performance decreases drastically for models with sequential token order in both *Rhyme-L* and *Rhyme-G*. With the same token order, the models with the " $\langle l_0 \rangle$ " of the local position aligned with the rhyming word perform slightly better than those not aligned. The results show that generating rhyming words before the other words in the sentences can significantly improve the rhyming performance of the generated outputs. However, the role of local position embeddings to help identify the rhyming words is less important. We hypothesize that, apart from the local position, the model can also identify the rhyming words from the global position (the tokens before or after the sentence delimiter tokens).

J Additional Examples

Generated Lyrics

闭上眼睛(qing)
Close your eyes,
携手向前行(xing)
and move forward hand in hand.
那黑夜的流星(xing)
the meteors of the night
燃烧生命(ming)
is burning tis life.

眼泪在不停(ing)
Tears are not stopping.
我愿为了你(ni)
For you,
凝聚成耀眼的流星(xing)
I am willing to condense into a dazzling meteor,
照在寂寞的每个夜里(li)
shining on every lonely night.
释放心中的光亮(liang)
Releasing the light in the heart
是迎向生命的光明(ming)
is to meet the light of the life.
尽情怀抱希望和梦想(xiang)
You can embrace your hopes and dreams,
痴情走到底(di)
sticking it until the end with infatuation.

闭上眼睛(qing)
Close your eyes,
携手向前行(xing)
and move forward hand in hand.
那黑夜的流星(xing)
the meteors of the night
燃烧生命(ming)
is burning tis life.

眼泪在不停(ing)
Tears are not stopping.
我愿为了你(ni)
For you,
凝聚成耀眼的流星(xing)
I am willing to condense into a dazzling meteor,
照在寂寞的每个夜里(li)
shining on every lonely night.
释放心中的光亮(liang)
Releasing the light in the heart
是迎向生命的光明(ming)
is to meet the light of the life.
尽情怀抱希望和梦想(xiang)
You can embrace your hopes and dreams,
痴情走到底(di)
sticking it until the end with infatuation.

献出你的天真和我的真心真意(yi)
Give your innocence and my sincerity.
你就是我的唯一(yi)
You are my only one.

凝聚成耀眼的流星(xing)
I am willing to condense into a dazzling meteor,
照在寂寞的每个夜里(li)
shining on every lonely night.
释放心中的光亮(liang)
Releasing the light in the heart
是迎向生命的光明(ming)
is to meet the light of the life.
尽情怀抱希望和梦想(xiang)
You can embrace your hopes and dreams,
痴情走到底(di)
sticking it until the end with infatuation.

Figure 5: An example generated by *SongRewriter*. End characters sharing the same vowel (the *pinyin* of the characters are in the adjacent bracket) are highlighted in the same color. Lyrics is splitted into blocks for clear illustration.

Generated Lyrics

谁是无尽夜迷人的烟花(hua)
Who is the charming fireworks of the endless night?
谁是谎言中美丽的童话(hua)
Who is the beautiful fairy tale among lies?
千言万语也无法来表达(da)
A thousand words can't express it.
谁会懂我对你牵挂(gua)
Who will understand my miss of you.

可是你懂我年少的年华(hua)
But you understand my youth.
你懂为爱我付出的代价(jia)
You know the price i pay for love.
千山万水也望尽了天涯(ya)
I look through thousands of mountains and rivers to the end of the world.
但愿能懂我对你牵挂(gua)
I wish you could understand my miss of you.

我已分不清谈情的真假(jia)
I can't tell the difference between true and false love.
也懂得如何如何去潇洒(sa)
But I know how to be cool.
我还厮守著心中的情话(hua)
I'm still holding on to the love words in my heart,
相伴每一个春夏秋冬(xia)
accompanying me every spring, autumn, winter and summer.

谁是无尽夜迷人的烟花(hua)
Who is the charming fireworks of the endless night?
谁是谎言中美丽的童话(hua)
Who is the beautiful fairy tale among lies?
千言万语也无法来表达(da)
A thousand words can't express it.
谁会懂我对你牵挂(gua)
Who will understand my miss of you.

你像是漫天刮起的风沙 sha)
You are like sand blowing from the sky.
我愿化作你身后的落霞(xia)
I would like to be the sunset behind you.
等待下一个温暖的盛夏(xia)
waiting for the next warm summer.
寒雨淋湿你一缕秀发 fa)
The cold rain wets a strand of your hair.

男人的眼泪也不愿流下(xia)
Men don't want to shed their tears.
我多想遮住心上的伤疤(ba)
I want to hide the scars on my heart.
想听你埋怨世界太复杂(za)
I want to hear you complain that the world is too complicated.
却学会了我对你牵挂(gua)
But I learned that I miss you.

我已分不清谈情的真假(jia)
I can't tell the difference between true and false love.
也懂得如何如何去潇洒(sa)
But I know how to be cool.
我还厮守著心中的情话(hua)
I'm still holding on to the love words in my heart,
相伴每一个春夏秋冬(xia)
accompanying me every spring, autumn, winter and summer.

泛黄了曾经走过的时光(guang)
The passed time has turned yellow.
它提醒了我思念很长(chang)
It reminds me that I have been missing you for a long time.
我的心已不再为你疯狂(kuang)
My heart is no longer crazy for you.
就让那深情被你遗忘(wang)
Just let that affectionate feeling be forgotten by you.

谁是无尽夜迷人的烟花(hua)
Who is the charming fireworks of the endless night?
谁是谎言中美丽的童话(hua)
Who is the beautiful fairy tale among lies?
千言万语也无法来表达(da)
A thousand words can't express it.
就让那深情被你融化(hua)
Just let that affection be melted by you

Figure 6: An example generated by *SongRewriter*. End characters sharing the same vowel (the *pinyin* of the characters are in the adjacent bracket) are highlighted in the same color. Lyrics is splitted into blocks for clear illustration.

Original Lyrics	Rhyme Scheme	Generated Lyrics
色白花青的锦鲤(i)跃然于碗底(d) <i>On the bottom of the white bowl, blue and white carps leap to and fro.</i>	-----i-----i	已没力气让自己(j)去把你忘记(j) <i>I don't have the strength to let myself forget you.</i>
临摹宋体(i)落款时却临记着你(ni) <i>Copying the Song inscription, to the end I am yet to sign, before you step into my mind.</i>	---i-----i	点点滴滴(d)已深深刻在脑海里(i) <i>Every little thing has been deeply imprinted in my mind.</i>
你隐藏在窑烧里(i)千年的秘密(mi) <i>The mystery of yours, buried deep in the kiln for thousands of years.</i>	-----i-----i	我不想再为了你(mi)而歇斯底里(i) <i>I don't want to be hysterical about you anymore.</i>
极细(ni)犹如绣花针落地(d) <i>is as delicate yet profound as a needle falling into the ground.</i>	---i-----i	怪自己(j)早已没有了勇气(qi) <i>Blame myself for not having the courage.</i>

Figure 7: An example of full song rewriting based on the internal rhyme scheme of the input lyrics. The lyrics is from *Blue and White Porcelain* by Jay Chou. The left column is the original lyrics, with the rhyme scheme in the middle column. For the rhyme scheme, only the vowels of the rhyming characters are remained. The vowels of other characters are replaced by a placeholder token. The output lyrics on the right column are generated based on the rhyme scheme.

Original Lyrics	Generated Lyrics
那时候的爱情 <i>The love then,</i>	年少时的爱情 <i>The love when you were young</i>
为什么就能那样简单 <i>why could it be so simple.</i>	是自以为的那样简单 <i>is as simple as you think.</i>
而又是为什么 <i>And why is it</i>	而又是为什么 <i>And why is that,</i>
人年少时 <i>that when people are young,</i>	年少轻狂 <i>being young and frivolous,</i>
一定要让深爱的人受伤 <i>they must hurt the person they deeply love.</i>	一定要让深爱的人受伤 <i>they must hurt the person they deeply love.</i>
在较相似的深夜里 <i>On this similar night,</i>	在那片青涩岁月里 <i>In those youthful years,</i>
你是否一样 <i>is it the same for you?</i>	你是否一样 <i>is it the same for you?</i>
也在静静追悔感伤 <i>Are you also quietly reminiscing the sadness?</i>	也在人前假装坚强 <i>Are you also pretending to be strong in front of people?</i>
如果当时我们能 <i>If we could</i>	如果当时我们能 <i>If we could</i>
不那么倔强 <i>have been less obstinate then,</i>	不那么倔强 <i>have been less obstinate then,</i>
现在也不那么遗憾 <i>we wouldn't be so regretful now.</i>	现在不必装模作样 <i>we don't need to pretend now.</i>

Figure 8: An example of partial sentence rewriting by *SongRewriter*. The original lyrics is from *Later* by Rene Liu. The inputs are the original lyrics with red tokens masked. The model generates tokens at the corresponding masked positions (highlighted in red on the right column).

Original Lyrics	Generated Lyrics
简单点 说话的方式简单点 <i>Keep it simple, keep the way you talk simple.</i>	简单点 说话的方式简单点 <i>Keep it simple, keep the way you talk simple.</i>
递进的情绪请省略 <i>Please omit the progressive emotions.</i>	递进的情绪请省略 <i>Please omit the progressive emotions.</i>
你又不是个演员 <i>You are not an actor.</i>	用你惯用的语言 <i>speak in your preferred language.</i>
别设计那些情节 <i>Don't design those plots.</i>	再排练好几千遍 <i>and rehearse thousands of times.</i>
没意见 我只想看看你怎么圆 <i>No opinion, I just want to see how you complete your lies.</i>	没意见 我只想看看你怎么圆 <i>No opinion, I just want to see how you complete your lies.</i>
你表演的太敷衍 <i>Your show of grief is too shallow.</i>	一番假作的敷衍 <i>a contrived grimace.</i>
像没天赋的演员 <i>like a talentless actor.</i>	几个虚假的语言 <i>a couple of lies.</i>
观众一眼能看见 <i>The audience can perceive at a glance.</i>	每一天都在上演 <i>are performed every day.</i>
该配合你演出的我 演视而不见 <i>I'm suppose to complement your acting but I ignored it.</i>	该配合你演出的我 演视而不见 <i>I'm suppose to complement your acting but I ignored it.</i>
在逼一个最爱你的人 即兴表演 <i>You're demanding the person who loves you the most to improvise.</i>	在逼一个最爱你的人 即兴表演 <i>You're demanding the person who loves you the most to improvise.</i>
什么时候我们开始 收起了底线 <i>When did we forget our bottom lines.</i>	什么时候我们开始 收起了底线 <i>When did we forget our bottom lines.</i>
顺应时代的改变 看那些拙劣的表演 <i>Watching those poor performances with the change of times.</i>	顺应时代的改变 看那些拙劣的表演 <i>Watching those poor performances with the change of times.</i>

Figure 9: An example of sentences rewriting by *SongRewriter*. The original lyrics is from *Actor* by Zhiqian Xue. The inputs are the original lyrics with sentences in red masked. The outputs are the red sentences on the right.

ACL 2023 Responsible NLP Checklist

A For every submission:

A1. Did you describe the limitations of your work?

6

A2. Did you discuss any potential risks of your work?

7

A3. Do the abstract and introduction summarize the paper's main claims?

1

A4. Have you used AI writing assistants when working on this paper?

Left blank.

B Did you use or create scientific artifacts?

Left blank.

B1. Did you cite the creators of artifacts you used?

Not applicable. Left blank.

B2. Did you discuss the license or terms for use and / or distribution of any artifacts?

Not applicable. Left blank.

B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?

Not applicable. Left blank.

B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?

our data does not contain any personal information

B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?

Not applicable. Left blank.

B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.

Appendix A

C Did you run computational experiments?

4

C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?

Appendix C

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Appendix C

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

4

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

4

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

4.2

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

4.3

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

Left blank.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

Left blank.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

Not applicable. Left blank.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

Not applicable. Left blank.