

SongGLM: Lyric-to-Melody Generation with 2D Alignment Encoding and Multi-Task Pre-Training

Jiaxing Yu¹, Xinda Wu¹, Yunfei Xu², Tiejiao Zhang¹, Songruoyao Wu¹, Le Ma¹, Kejun Zhang^{1,3*}

¹College of Computer Science and Technology, Zhejiang University

²AI Center, Guangdong OPPO Mobile Telecommunications Corp., Ltd.

³Innovation Center of Yangtze River Delta, Zhejiang University

{yujx, wuxinda}@zju.edu.cn, xuyunfei@oppo.com, {kretzer0421, wsry, maller, zhangkejun}@zju.edu.cn

Abstract

Lyric-to-melody generation aims to automatically create melodies based on given lyrics, requiring the capture of complex and subtle correlations between them. However, previous works usually suffer from two main challenges: 1) lyric-melody alignment modeling, which is often simplified to one-syllable/word-to-one-note alignment, while others have the problem of low alignment accuracy; 2) lyric-melody harmony modeling, which usually relies heavily on intermediates or strict rules, limiting model’s capabilities and generative diversity. In this paper, we propose SongGLM, a lyric-to-melody generation system that leverages 2D alignment encoding and multi-task pre-training based on the General Language Model (GLM) to guarantee the alignment and harmony between lyrics and melodies. Specifically, 1) we introduce a unified symbolic song representation for lyrics and melodies with word-level and phrase-level (2D) alignment encoding to capture the lyric-melody alignment; 2) we design a multi-task pre-training framework with hierarchical blank infilling objectives (n-gram, phrase, and long span), and incorporate lyric-melody relationships into the extraction of harmonized n-grams to ensure the lyric-melody harmony. We also construct a large-scale lyric-melody paired dataset comprising over 200,000 English song pieces for pre-training and fine-tuning. The objective and subjective results indicate that SongGLM can generate melodies from lyrics with significant improvements in both alignment and harmony, outperforming all the previous baseline methods.

1 Introduction

Lyric-to-melody generation, which aims to automatically generate melodies from given lyrics, has attracted lots of attention from both academia and industry. When creating melodies, capturing the complex and subtle lyric-melody correlations is crucial. Previous works (Watanabe et al. 2018; Bao et al. 2019; Yu, Srivastava, and Canales 2021; Ju et al. 2022; Sheng et al. 2021; Lv et al. 2022; Zhang et al. 2022; Ding et al. 2024; Wang et al. 2024; Yu et al. 2024) in this field have achieved great progress in capturing these correlations, but still encounter two primary challenges: lyric-melody alignment modeling and harmony modeling.

*Corresponding author.




Figure 1: Illustration of lyric-to-melody generation challenges: alignment modeling and harmony modeling.

1) Lyric-melody alignment modeling. Lyric-melody alignment denotes the quantitative relationships between syllables/words and notes, particularly the number of notes mapped to a single syllable or word, which has a significant impact on the richness and singability of a song. Since most existing works (Yu, Srivastava, and Canales 2021; Ju et al. 2022; Lv et al. 2022; Zhang et al. 2022) only explore the one-syllable/word-to-one-note (one-to-one) alignment, few works consider the one-syllable/word-to-multiple-notes (one-to-multiple) alignment. Bao et al. (2019) predicted the number of notes corresponding to the given syllable with a greedy alignment strategy. Sheng et al. (2021) utilized sentence-level and token-level attention masks to achieve alignment between word/syllable and note. However, these methods usually suffer from low alignment accuracy, attributable to their indirect ways of learning the lyric-melody alignment. Consequently, it is essential to introduce a unified representation for lyrics and melodies that directly captures the lyric-melody alignment.

2) Lyric-melody harmony modeling. Lyric-melody harmony refers to the qualitative relationships between syllables/words and notes, emphasizing their feature coherence, which is crucial for the rhythmic and structural consistency of a song. Ju et al. (2022) proposed templates which consist of tonality, chord, rhythm, and cadence, serving as a bridge between lyrics and melodies to improve the harmony. Lv et al. (2022) extracted key features from lyrics, including tonality, rhythm, chord, and structure, and leveraged these features as the query to retrieve and concatenate pre-generated melody segments. Zhang et al. (2022) introduced several rules about lyric-melody relationships from the perspectives of tone, rhythm, and structure, and integrated them into the decoding step of lyric-to-melody generation models. These approaches rely heavily on either intermediates (tem-

plates or keys) or strict rules which limit model’s capabilities and generative diversity. Therefore, proposing a method that ensures the lyric-melody harmony while maintaining generative creativity is essential.

Specifically, to address the challenge of lyric-melody alignment modeling, we introduce a unified symbolic song representation that provides a comprehensive method to encode lyric and melodic information. For both words in lyrics and notes in melodies, the representation consists of three types of attributes: generic, content-related, and alignment-related. Generic attributes refer to basic properties that apply across both words and notes, such as token types. Content-related attributes contain distinct elements describing words or notes, such as the textual contents of words and the musical features of notes. Alignment-related attributes include word and phrase level alignment ids (word ID and phrase ID), serving as 2D alignment encoding that directly provide hierarchical alignment information between lyrics and melodies.

To handle the challenge of lyric-melody harmony modeling, we propose a multi-task pre-training framework based on GLM (Du et al. 2022) that enables the model to capture multi-scale, multi-dimensional harmony between lyrics and melodies. We concatenate the word sequence from the lyrics with the note sequence from the melody, using the word sequence as a condition. Then, we create hierarchical blank infilling objectives (n-gram, phrase, and long span) from the perspective of word, phrase and song, to jointly pre-train the model by blanking out continuous spans of tokens from the note sequence and contextually reconstructing these spans. Since different n-grams contribute differently to the lyric-melody harmony, we explore the interaction of lyric and melodic features, including syllable stress, melodic peak and rhythm skeleton, and introduce two lyric-melody relationships between them. By incorporating these relationships into the process of n-gram extraction, we can select harmonized n-grams that best represent the significant and repeating patterns in the lyric-melody harmony. Furthermore, we construct a large-scale lyric-melody paired dataset based on MelodyNet (Wu et al. 2023), that contains more than 200,000 English song pieces for pre-training and fine-tuning.

Our contributions are summarized as follows: 1) We propose SongGLM, a lyric-to-melody generation system that effectively tackles the challenges of lyric-melody alignment and harmony modeling. 2) We design 2D alignment encoding through a unified symbolic song representation to ensure the alignment between generated melodies and corresponding lyrics. 3) We propose a multi-task pre-training framework based on GLM with hierarchical blank infilling objectives (n-gram, phrase, and long span), and incorporate lyric-melody relationships into the extraction of harmonized n-grams to improve the harmony between lyrics and melodies. 4) Objective and subjective evaluation results show that SongGLM can generate high-quality melodies from lyrics with significant improvements in both alignment and harmony, outperforming all the previous baseline methods. This highlights the effectiveness of 2D alignment encoding and multi-task pre-training in lyric-to-melody generation.

2 Background

2.1 Lyric-to-Melody Generation

Over the past few years, there have been advancements in deep learning approaches for lyric-to-melody generation. Bao et al. (2019), Lee, Fang, and Ma (2019), and Yu, Srivastava, and Canales (2021) adopted end-to-end models to generate melodies from lyrics. These methods cannot fully capture the relationships between lyrics and melodies due to the limited availability of paired lyric-melody dataset. To address this issue, Ju et al. (2022) divided the generation process into two stages, lyric-to-template and template-to-melody, to leverage unpaired data. Lv et al. (2022) proposed a generation-retrieval pipeline by sharing same key features between lyrics and melodies. Sheng et al. (2021) trained lyrics generation and melody generation models separately with unpaired data and performed attention based alignment modeling. Zhang et al. (2022) developed an expert system on Sheng et al. (2021) and Ju et al. (2022) that incorporated lyric-melody relationships from the music theory to improve lyric-melody harmony. However, the aforementioned studies are inadequate in effectively handling the complex and subtle correlations between lyrics and melodies, particularly failing to address both alignment modeling and harmony modeling concurrently. In this paper, we propose SongGLM, a novel lyric-to-melody generation system with 2D alignment encoding and multi-task pre-training to tackle these challenges.

2.2 Pre-Training Frameworks

Pre-training frameworks have made significant contributions to the development of automatic music composition. Early encoder-only frameworks, like BERT (Devlin et al. 2019), adopted multi-layer bidirectional Transformer encoders to learn deep bidirectional representations, showing strong capabilities in music understanding tasks (Wang and Xia 2021; Zeng et al. 2021). Later decoder-only frameworks, like GPT (Radford et al. 2018), leveraged multi-layer unidirectional Transformer decoders to capture rich context information, which are suited for music generation tasks (Ens and Pasquier 2020). Encoder-decoder frameworks, like MASS (Song et al. 2019), integrated encoder and decoder modules for better understanding and generating music sequences (Sheng et al. 2021). Recently, more and more innovative and efficient frameworks (Dong et al. 2019; Raffel et al. 2020; Du et al. 2022; Touvron et al. 2023) have been proposed and adopted in different downstream tasks. Among them, GLM (Du et al. 2022), based on a unified encoder-decoder framework with autoregressive blank infilling, has shown promising results in music generation tasks (Wu et al. 2023). In this paper, we exploit the powerful capabilities of GLM, and build SongGLM to address the challenges encountered in lyric-to-melody generation.

3 Method

3.1 System Overview

An overview of SongGLM is shown in Figure 2. Given the paired lyric-melody dataset, we first establish two relation-

ships between lyrics and melodies based on their representative features, and incorporate these relationships into n-gram extraction to select the most harmonized n-grams. Then, we introduce a unified symbolic song representation with 2D alignment encoding and adopt a multi-task pre-training framework that employs hierarchical blank infilling objectives for lyric-to-melody generation. In the following subsections, we describe the details of harmonized n-gram extraction and lyric-to-melody generation.

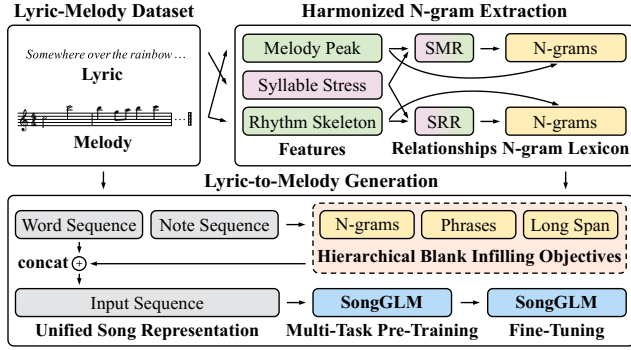


Figure 2: An overview of SongGLM. It includes two stages: harmonized n-gram extraction (detailed in Section 3.2) and lyric-to-melody generation (detailed in Section 3.3).

3.2 Harmonized N-gram Extraction

N-grams are widely used in Music Information Retrieval for understanding and generating tasks (Zheng, Moh, and Moh 2017; Xiao et al. 2021; Wu et al. 2023). However, existing n-gram extraction methods (Zeng et al. 2021; Wu et al. 2023) often struggle to effectively capture the harmony between lyrics and melodies when applied to lyric-to-melody generation tasks. To address this challenge, we first introduce the three most representative features of lyrics and melodies, and establish the qualitative relationships between them based on these features. Then, we propose a novel n-gram extraction strategy that incorporates these relationships to select the most harmonized n-grams.

Features

We describe three key features of lyrics and melodies and the corresponding extraction functions (f_w or f_n) across the dimension of word, pitch and time, specifically: **syllable stress**, **melodic peak**, and **rhythm skeleton**.

Syllable Stress. Syllable stress (Ladefoged, Draper, and Whitteridge 1958) refers to the emphasis placed on particular syllables within words. The sequence of syllable stress represents the rhythmic pattern of lyrics, and plays a crucial role in lyric-to-melody generation. According to the CMU Pronouncing Dictionary¹, each syllable stress can be categorized into three levels on an ordinal scale: Unstressed, Primary Stress, and Secondary Stress. For the ‘‘Syllable Stress’’ feature, we define f_w as follows:

$$f_w(W) = [s_1, \dots, s_y] \quad (1)$$

¹<https://speech.cs.cmu.edu/cgi-bin/cmudict>

where $s_i \in \{0, 1, 2\}$ represents the stress level of the i^{th} syllable in a word, with 0 indicating Unstressed, 1 indicating Primary Stress, and 2 indicating Secondary Stress.

Melodic Peak. Melodic peaks (Eitan 2016) refer to the notes with a higher pitch compared to the preceding and subsequent notes. The sequence of melodic peaks describes the movement pattern of the melody among high pitches. For the ‘‘Melodic Peak’’ feature, given the pitch sequence $P = [p_1, \dots, p_n]$ of the melody, we define $f_{n.MP}$ for the i^{th} note as follows:

$$f_{n.MP}(N) = \begin{cases} 1, & 1 < i < n, p^i > p^{i-1}, p^i > p^{i+1} \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

Rhythm Skeleton. The rhythm skeleton (Zhang et al. 2023) represents a set of specific notes that are acoustically more prominent than others, due to the joint effect of meter and rhythm on the time dimension. Following (Zhang et al. 2023), we extract metrical accents, agogic accents on metrical accents, and agogic accents on syncopations as the rhythm skeleton. For the ‘‘Rhythm Skeleton’’ feature, given the rhythm skeleton sequence $RS = [N_1, \dots, N_z]$ of the melody, we define $f_{n.RS}$ for the i^{th} note as follows:

$$f_{n.RS}(N) = \begin{cases} 1, & N_i \in RS \\ 0, & N_i \notin RS \end{cases} \quad (3)$$

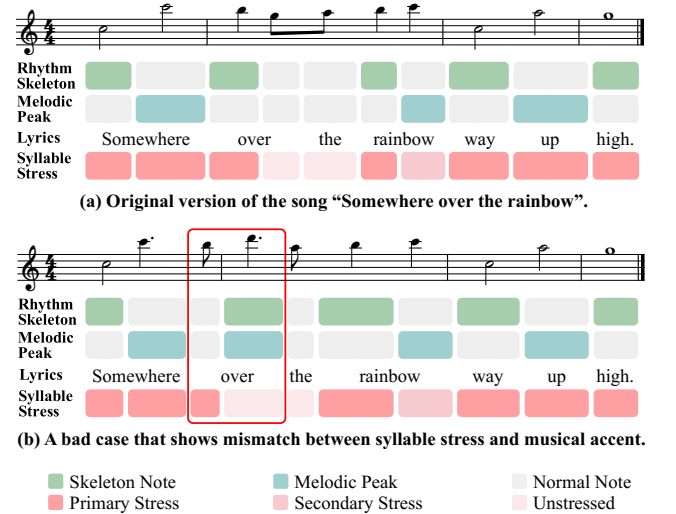


Figure 3: An example of lyric and melodic features, as well as the relationships between them.

Relationships

Since stressed syllables are often associated with musically accented notes (Nichols et al. 2009), we reveal the mechanism of interaction between lyric and melodic features and introduce two relationships.

Syllable Stress and Melodic Peak (SMR). Composers usually employ melodic peaks to emphasize certain syllables in songwriting. Specifically, as shown in Figure 3(a), a

high level of syllable stress tends to occur with the melodic peak.

Syllable Stress and Rhythm Skeleton (SRR). Rhythm skeleton is another method by which composers highlight specific words in a song. For instance, as illustrated in Figure 3(a), syllables with high stress levels are often associated with the note in the rhythm skeleton.

The correspondence between syllable stress and melodic accents can significantly improve the harmony between lyrics and melodies. On the other hand, if a melodic accent mismatches with syllable stress (red block in Figure 3(b)), it may disrupt the natural flow of the song, potentially resulting in disharmony that can detract from the overall listening experience.

Extraction Strategy

To capture the above relationships and ensure harmony between lyrics and melodies, we propose a novel n-gram extraction strategy. This strategy involves calculating a composite score for each n-gram, which includes both a melodic score and a lyric-melody relationship score. N-grams with high composite scores are selected as harmonized n-grams. The details of this strategy are outlined below.

Given a word sequence W of lyrics, a note sequence N of melody, and paired feature extraction functions f_w and f_n ($f_{n.MP}$ or $f_{n.RS}$) from lyric-melody relationships, we denote each word-note feature pair as a joint uni-gram $\{(f_w(w), f_n(N_w)) \mid w \in W, N_w \subseteq N\}$, where N_w represents the set of notes corresponding to the single word w . Subsequently, we extract joint n-grams for n ranging from 2 to 12, comprising lyric n-grams $f_w(W_n)$ with n words and melodic n-grams $f_n(N_{W_n})$. Furthermore, we compute t-statistic scores s_t (Xiao et al. 2021) for the lyric and melodic n-grams separately, s_l and s_m . Each melodic n-gram $f_n(N_{W_n})$ is associated with a set of m distinct lyric n-grams $F_w(W_n^m) = \{f_w(W_n^1), \dots, f_w(W_n^m)\}$ from different joint n-grams. Finally, the score s of a joint n-gram $\{(f_w(W_n), f_n(N_{W_n})) \mid W_n \subseteq W, N_{W_n} \subseteq N\}$ consists of two parts (the melodic n-gram t-statistic score s_m and the lyric-melody relationship score s_{lm}), which is defined as:

$$s = s_m + s_{lm} \quad (4)$$

$$s_l = s_t(f_w(W_n)) \quad (5)$$

$$s_m = s_t(f_n(N_{W_n})) \quad (6)$$

$$s_{lm} = C(F_w(W_n^m)) \cdot \frac{1}{m} \sum_{i=1}^m s_l^i \quad (7)$$

$$C(F_w(W_n^m)) = \begin{cases} 1, & m = 1 \\ 1 - H'(F_w(W_n^m)), & m > 1 \end{cases} \quad (8)$$

$$H'(F_w(W_n^m)) = \frac{-\sum_{i=1}^m p(W_n^i) \log p(W_n^i)}{\log m} \quad (9)$$

where p represents the occurrence probability of a given lyric n-gram among all corresponded lyric n-grams to the

melodic n-gram, C represents the concentration of the lyric n-gram set associated with the melodic n-gram, derived from the normalized entropy H' . The higher the concentration C , the better the joint n-gram represents a significant and repeating pattern in the lyric-melody relationship, thereby more effectively influencing the harmony between lyrics and melodies. Based on their scores, we select the top 25% of joint n-grams as harmonized n-grams to construct the final n-gram lexicon for word-level sampling in lyric-to-melody generation.

3.3 Lyric-to-Melody Generation

On top of the above extracted harmonized n-grams, we build SongGLM upon GLM (Du et al. 2022) with a single Transformer-based encoder-decoder framework for lyric-to-melody generation, as shown in Figure 4. In the pre-training stage, we adopt a multi-task pre-training framework with hierarchical blank infilling objectives. In the fine-tuning and inference stage, we utilize causal language modeling to predict the next note sequentially from left to right.

Unified Symbolic Song Representation

Inspired by OctupleMIDI (Zeng et al. 2021), we design a unified symbolic song representation for lyric-to-melody generation that allows the model to learn the lyric-melody alignment in an efficient and direct way. It consists of three different types of tokens: **Word**, **Note** and **Special**, each containing three sets of attributes: content-related, alignment-related, and generic. We consolidate the attributes into a single compound token to reduce the sequence length.

For **Word** and **Note** tokens, we assign the same alignment-related and generic attributes but different content-related attributes. Specifically, alignment-related attributes include two alignment ids, as shown in Figure 4(a). The first alignment id represents word-level alignment, called Word ID. For each **Word** token, it denotes the position in the word sequence, starting from 0 to the total length - 1. For each **Note** token, it equals to the Word ID of the word corresponding to the note. The second alignment id represents phrase-level alignment, named Phrase ID. For both tokens, the Phrase IDs refer to the musical phrase to which they belong. The above two alignment ids are encoded into embedding vectors, serving as **2D alignment encoding** to guarantee the hierarchical alignments between words and notes. Generic attributes include token types, which enhance the model’s capacity to differentiate between **Word** and **Note**. And for content-related attributes, **Note** tokens comprise five musical elements: bar, position, pitch, duration, and tempo, while **Word** tokens contain the text of the word. We only select words that are included in the CMU Pronouncing Dictionary and sorted them according to their frequency of occurrence in the lyrics. To facilitate computational modeling, we set content-related attributes of the **Word** token to *None* in the **Note** token, and vice versa.

For **Special** tokens, we adopt five special delimiter symbols: $\langle BOS \rangle$, $\langle EOS \rangle$, $\langle MASK \rangle$, $\langle PAD \rangle$, and $\langle SEP \rangle$. Similar to the **Word** and **Note** tokens, every **Special** token contains all attributes, each bearing the same value as itself.

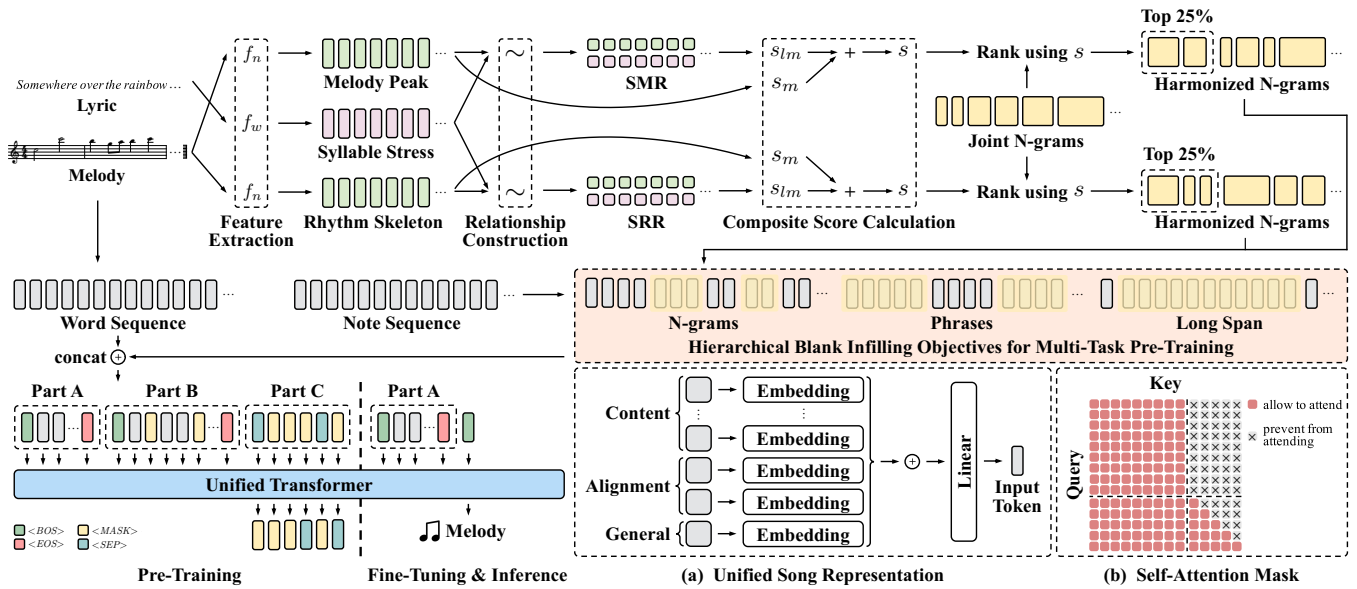


Figure 4: Illustration of our proposed SongGLM. In the harmonized n-gram extraction stage (top), we construct two lyric-melody relationships based on their features, and incorporate these relationships into n-gram extraction to select the top 25% as harmonized n-grams. In the lyric-to-melody generation stage (bottom), we leverage the extracted n-grams and propose three hierarchical blank infilling objectives for multi-task pre-training. We adopt casual language modeling for fine-tuning and inference. (a) Unified song representation for both lyrics and melodies. (b) Self-attention mask in the unified transformer.

Multi-Task Pre-Training

Multi-task pre-training has been shown to enhance model’s performance in a variety of tasks (Sun et al. 2021; Wu et al. 2023). Meanwhile, autoregressive blank infilling is an effective pre-training approach for language models (Du et al. 2022; Wu et al. 2023). Following their success, we implement a multi-task pre-training framework with hierarchical autoregressive blank infilling objectives in SongGLM.

Autoregressive blank infilling involves blanking out continuous spans of tokens from the input sequence and contextually reconstructing these spans during model training. Given an input sequence $S = [W_1, \dots, W_m, N_1, \dots, N_n]$, multiple token spans $s = \{s_1, \dots, s_k\}$ are sampled from the note sequence N . Each span s_i corresponds to a series of consecutive tokens $[s_{i,1}, \dots, s_{i,l_i}]$ in N , and is replaced with a single special token $\langle \text{MASK} \rangle$, forming a corrupted token sequence S_{corrupt} . The model is trained to predict the missing tokens within the spans from the corrupted token sequence in an autoregressive way, with access to the corrupted token sequence and previously predicted spans.

We construct three hierarchical autoregressive blank infilling objectives for pre-training to capture the multi-scale, multi-dimensional harmony between lyrics and melodies.

Word-Level. Based on the extracted n-gram lexicon, we randomly sample two types of harmonized n-grams from the note sequence with the Maximum Matching Algorithm (Xiao et al. 2021). The total length of sampled n-grams constitutes 15% of the note sequence. We replace each sampled n-gram with 1) the $\langle \text{MASK} \rangle$ token 80% of the time, 2) a random n-gram 10% of the time, and 3) the original n-gram 10% of the time. These objectives aims to capture word-note

level harmony between lyrics and melodies.

Phrase-Level. Multiple musical phrases are sampled from the note sequence, with the total length accounting for 50% of the original note sequence length. We consider both lyric and melodic information for musical phrase boundary recognition. This objective aims to capture lyric-phrase level harmony between lyrics and melodies, as well as to ensure the coherence of melodic contexts.

Song-Level. We sample a single long span that covers 50% of the original note tokens. This objective aims to improve the overall harmony between lyrics and melodies, and enhance the model’s ability of melodic structure modeling.

Sequence Modeling

Pre-Training. In the pre-training stage, the input sequence S contains three parts: Part A is the word sequence W , Part B is the corrupted note sequence, and Part C consists of the masked spans with each separated by a $\langle \text{SEP} \rangle$ token. Tokens in Part A & Part B form the corrupted sequence S_{corrupt} , and can attend to each other. Part C tokens can only attend to preceding tokens, and tokens in Part A & Part B. Figure 4(b) illustrates how attention weight is modified through the self-attention mask to control the token’s attention. With this mechanism, our unified model effectively learns a bidirectional encoder for Part A & Part B, and a unidirectional decoder for Part C.

Fine-Tuning and Inference. In the fine-tuning and inference stage, we employ causal language modeling. The input sequence begins with the word sequence (Part A) and a $\langle \text{BOS} \rangle$ token (indicating the start of the note sequence), and the model predicts the next token in an autoregressive manner until it generates an $\langle \text{EOS} \rangle$ token.

4 Experiments

4.1 Lyric-Melody Dataset

A large-scale paired dataset is critical for lyric-to-melody generation models to capture lyric-melody correlations and attain superior performance. However, the current largest paired dataset (Yu, Srivastava, and Canales 2021) only contains 12,197 MIDI songs and lacks one-to-multiple alignment. In this paper, we acquire approximately 1.6 million raw MIDI data from MelodyNet (Wu et al. 2023), and construct a large-scale lyric-melody paired dataset with varied word-note alignments, including both one-to-one and one-to-multiple alignments.

To obtain high-quality MIDI song pieces from the raw MIDI data, we perform data processing in four phases: lyric processing phase, melody processing phase, lyric-melody combined processing phase, and de-duplication phase. After data processing, the final dataset contains 206,884 English MIDI song pieces with 4,921.79 hours of melodies in total. We extract LMD-full dataset and Reddit-sourced dataset from the processed dataset, with a total of 8,195 pieces, for fine-tuning and use the remaining part for pre-training.

4.2 Model Configuration

SongGLM uses a single Transformer (Vaswani et al. 2017) as the basic model structure, and is pre-trained in two versions: 1) SongGLM_{small} on the small-scale dataset, containing 40,000 songs randomly selected from the full pre-training dataset, which aims to compare with baseline models that are also pre-trained on the small-scale dataset; 2) SongGLM_{base} on the full pre-training dataset, for demonstrating the best capability of SongGLM and presenting the state-of-the-art results. We adopt our proposed multi-task pre-training framework for both SongGLM_{small} and SongGLM_{base}.

4.3 Evaluation Metrics

In this subsection, we briefly introduce the objective and subjective metrics used to evaluate the performance of SongGLM for lyric-to-melody generation.

Objective Metrics. We consider four widely used objective metrics (Sheng et al. 2021; Ju et al. 2022; Zhang et al. 2022) to evaluate lyric-melody harmony by measuring the similarity between generated and ground-truth melodies, in terms of pitch (D_P), duration (D_D), IOI (D_{IOI}) (Yang and Lerch 2020), and the overall melody sequence (MD). Besides, we propose alignment distribution similarity (D_A) to measure the consistency of word-note alignments across the generated and ground-truth melodies. To reduce the effects of stochastic sampling, we conduct each experiment on the test set 10 times.

Subjective Metrics. For subjective evaluation, we conduct a human listening test and compare SongGLM_{small} with the original and ReLyMe-equipped SongMASS and TeleMelody. We apply each model to generate 15 samples randomly and invite 10 participants, where 6 of them can understand basic music theory, to evaluate these samples. Specifically, we require all participants to score each sample using a ten-point scale (1 for lowest and 10 for highest) from

two aspects: 1) the quality of the generated melody; 2) the quality of the overall sample, considering both the melody and corresponding lyrics. We utilize ACE Studio² to synthesize the singing voice from lyrics and melodies, and provide participants with the lyrics, melodies and singing voice.

4.4 Main Results

To verify the effectiveness of SongGLM in the alignment and harmony between lyrics and melody, we compare SongGLM_{small} to the original and ReLyMe-equipped SongMASS and TeleMelody with same system configurations. The objective results are shown in Table 1. It is evident that SongGLM_{small} significantly surpasses the baseline models across all objective metrics. Specifically, D_A indicates that SongGLM_{small} outperforms in lyric-melody alignment, while D_P , D_D , D_{IOI} and MD suggest that SongGLM_{small} is the most capable of ensuring the harmony between lyrics and melodies. Table 2 shows the subjective results, from which we can see that for melody itself, SongGLM_{small} can generate diverse and consistent melodies. For the overall song, SongGLM_{small} not only ensures the rhythmic and structural consistency between lyrics and melody, but also achieves the best results in singability and overall performance. Besides, SongGLM_{base} achieves better results with a larger model and pre-training dataset in both objective and subjective evaluations, showing the scalability and capability of SongGLM.

4.5 Method Analysis

In this subsection, we analyze the effects of the designed 2D alignment encoding, lyric-melody relationships, and multi-task pre-training by conducting experiments on the following nine settings of SongGLM_{small} with same configurations. Table 1 presents the overall objective results of these settings.

- **Scratch:** train from scratch on the full dataset, without pre-training nor 2D alignment encoding.
- **CLM:** pre-train using casual language modeling (CLM).
- **CLM – 2D Alignment Encoding:** pre-train using CLM, without 2D alignment encoding.
- **Random:** pre-train using a random span sampling strategy (Joshi et al. 2020).
- **Harmonized N-gram:** pre-train using our proposed harmonized n-gram sampling strategy.
- **Harmonized N-gram – SMR:** pre-train using our proposed harmonized n-gram sampling strategy, without syllable stress and melodic peak relationship (SMR).
- **Harmonized N-gram – SRR:** pre-train using our proposed harmonized n-gram sampling strategy, without syllable stress and rhythm skeleton relationship (SRR).
- **Phrase:** pre-train using a random phrase sampling strategy.
- **Long:** pre-train using a single long span sampling strategy (Du et al. 2022).

²<https://www.acestudio.ai>

Model	Alignment		Harmony		
	$D_A(\%) \uparrow$	$D_P(\%) \uparrow$	$D_D(\%) \uparrow$	$D_{IOI}(\%) \uparrow$	$MD \downarrow$
SongMASS	-	87.25 ± 1.39	75.79 ± 1.23	81.94 ± 2.02	8.48 ± 0.75
TeleMelody	-	89.62 ± 1.12	84.55 ± 1.52	79.38 ± 0.88	6.36 ± 0.89
ReLyMe (in SongMASS)	-	90.25 ± 0.80	84.65 ± 0.93	86.69 ± 1.04	6.98 ± 0.80
ReLyMe (in TeleMelody)	-	92.19 ± 0.72	87.52 ± 1.02	84.80 ± 1.17	5.90 ± 0.88
SongGLM _{small}	94.32 ± 0.64	96.48 ± 0.94	95.34 ± 0.99	93.44 ± 0.72	4.17 ± 0.21
SongGLM _{base}	96.83 ± 0.59	96.50 ± 0.71	96.48 ± 0.97	94.10 ± 0.93	3.85 ± 0.30
Scratch	84.15 ± 0.69	88.61 ± 0.83	86.97 ± 1.09	82.91 ± 1.03	6.29 ± 0.84
CLM	90.20 ± 0.66	91.90 ± 0.68	90.69 ± 1.04	88.09 ± 0.99	4.95 ± 0.32
– 2D Alignment Encoding	83.57 ± 0.89	92.12 ± 0.84	87.40 ± 1.02	84.61 ± 1.11	5.67 ± 0.60
Random	93.87 ± 0.68	92.84 ± 0.65	92.06 ± 1.11	89.07 ± 1.21	4.68 ± 0.52
Harmonized N-gram	93.74 ± 0.51	94.11 ± 0.67	93.65 ± 1.09	91.88 ± 1.05	4.40 ± 0.48
– SMR	93.67 ± 0.60	91.77 ± 0.78	92.98 ± 1.06	91.13 ± 1.03	4.79 ± 0.41
– SRR	93.54 ± 0.59	93.93 ± 0.68	90.81 ± 1.03	89.04 ± 1.10	4.86 ± 0.45
Phrase	92.95 ± 0.58	94.09 ± 0.66	93.71 ± 1.21	91.72 ± 1.17	4.66 ± 0.25
Long	93.54 ± 0.67	94.24 ± 0.63	93.67 ± 1.06	91.68 ± 1.10	4.60 ± 0.49

Table 1: Objective results of SongGLM with different settings and baseline systems (Mean ± SD). SMR refers to syllable stress and melodic peak relationship, and SRR refers to syllable stress and rhythm skeleton relationship.

Model	Melody			Melody + Lyrics			Overall
	Richness	Consistency	Listenability	Rhythmicity	Structure	Singability	
SongMASS	6.10 ± 0.39	6.03 ± 0.36	5.91 ± 0.40	5.76 ± 0.39	5.78 ± 0.28	5.68 ± 0.38	5.77 ± 0.32
TeleMelody	6.54 ± 0.36	6.38 ± 0.35	6.39 ± 0.42	6.33 ± 0.42	6.26 ± 0.38	6.38 ± 0.42	6.45 ± 0.28
ReLyMe (in SongMASS)	6.57 ± 0.38	6.33 ± 0.38	6.39 ± 0.38	6.46 ± 0.44	6.41 ± 0.34	6.39 ± 0.45	6.51 ± 0.37
ReLyMe (in TeleMelody)	7.12 ± 0.30	7.01 ± 0.36	7.06 ± 0.36	6.91 ± 0.30	6.93 ± 0.27	6.81 ± 0.35	6.95 ± 0.31
SongGLM _{small}	7.28 ± 0.32	7.47 ± 0.29	7.41 ± 0.41	7.48 ± 0.33	7.63 ± 0.27	7.57 ± 0.47	7.60 ± 0.37
SongGLM _{base}	7.54 ± 0.28	7.78 ± 0.38	7.74 ± 0.34	7.66 ± 0.25	7.75 ± 0.28	7.83 ± 0.39	7.79 ± 0.31

Table 2: Subjective results of SongGLM and baseline systems. Each score is calculated with 95% confidence intervals.

Effectiveness of 2D Alignment Encoding. To verify the effectiveness of 2D alignment encoding, we compare the performance of CLM-based SongGLM_{small} with and without 2D alignment encoding. For settings without 2D alignment encoding, we assign notes to each word in lyrics, ensuring that the number of notes for each word equals to the number of vowels in the word. As shown in Table 1, our proposed 2D alignment encoding achieve much better scores on D_A , due to its excellent ability to directly capture the alignment between lyrics and melodies.

Effectiveness of Lyric-Melody Relationships. To verify the contribution of each relationship, we conduct experiments based on the harmonized n-gram setting, excluding each relationship separately. The results in Table 1 show that the syllable stress and melodic peak relationship mainly contributes to pitch harmony (D_P) between lyrics and melodies, while the syllable stress and rhythm skeleton relationship plays an important role in rhythm harmony (D_{IOI} and D_D).

Effectiveness of Multi-Task Pre-Training. To further explore the benefits of multi-task pre-training, we compare it with single-task pre-training: 1) Harmonized N-gram, 2) Phrase, and 3) Long. The results in Table 1 show that our proposed multi-task pre-training method achieves the highest performance among all settings, demonstrating its excel-

lent modeling performance on lyric-to-melody generation.

5 Conclusion

In this paper, we propose SongGLM, a lyric-to-melody generation system that leverages 2D alignment encoding and multi-task pre-training to ensure the alignment and harmony between lyrics and melodies. We introduce a unified symbolic song representation for lyrics and melodies that contains generic, content-related, and alignment-related attributes, and 2D alignment encoding to capture accurate alignments between lyrics and melodies. We design a multi-task pre-training framework with hierarchical blank infilling objectives (n-gram, phrase, and long span), and integrate lyric-melody relationships into the extraction of harmonized n-grams to guarantee the harmony between lyrics and melodies. Both objective and subjective results indicate that our proposed SongGLM can generate high-quality melodies from lyrics with remarkable lyric-melody alignment and harmony. Furthermore, method analysis shows the effectiveness of the detailed designs in SongGLM. In the future, we plan to extend our research to include more languages, such as Chinese, and explore the application of SongGLM to other automatic music composition tasks, such as text-to-music generation and video-to-music generation.

Acknowledgments

This research was supported by the National Key Research and Development Program of China (No. 2023YFF0904900), the National Natural Science Foundation of China (No. 62272409), Guangdong OPPO Mobile Telecommunications Corp., Ltd., and ACE Studio.

References

- Bao, H.; Huang, S.; Wei, F.; Cui, L.; Wu, Y.; Tan, C.; Piao, S.; and Zhou, M. 2019. Neural melody composition from lyrics. In *CCF International Conference on Natural Language Processing and Chinese Computing*, 499–511. Springer.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 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)*, 4171–4186. Association for Computational Linguistics.
- Ding, S.; Liu, Z.; Dong, X.; Zhang, P.; Qian, R.; He, C.; Lin, D.; and Wang, J. 2024. Songcomposer: A large language model for lyric and melody composition in song generation. *arXiv preprint arXiv:2402.17645*.
- Dong, L.; Yang, N.; Wang, W.; Wei, F.; Liu, X.; Wang, Y.; Gao, J.; Zhou, M.; and Hon, H.-W. 2019. Unified language model pre-training for natural language understanding and generation. *Advances in neural information processing systems*, 32.
- Du, Z.; Qian, Y.; Liu, X.; Ding, M.; Qiu, J.; Yang, Z.; and Tang, J. 2022. Glm: General language model pretraining with autoregressive blank infilling. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 320–335. Association for Computational Linguistics.
- Eitan, Z. 2016. *Highpoints: A study of melodic peaks*. University of Pennsylvania Press.
- Ens, J.; and Pasquier, P. 2020. Mmm: Exploring conditional multi-track music generation with the transformer. *arXiv preprint arXiv:2008.06048*.
- Joshi, M.; Chen, D.; Liu, Y.; Weld, D. S.; Zettlemoyer, L.; and Levy, O. 2020. Spanbert: Improving pre-training by representing and predicting spans. *Transactions of the association for computational linguistics*, 8: 64–77.
- Ju, Z.; Lu, P.; Tan, X.; Wang, R.; Zhang, C.; Wu, S.; Zhang, K.; Li, X.; Qin, T.; and Liu, T.-Y. 2022. TeleMelody: Lyric-to-Melody Generation with a Template-Based Two-Stage Method. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 5426–5437. Association for Computational Linguistics.
- Ladefoged, P.; Draper, M. H.; and Whitteridge, D. 1958. Syllables and stress. *Le Maître Phonétique*, 36: 1–14.
- Lee, H.-P.; Fang, J.-S.; and Ma, W.-Y. 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)*, 84–88.
- Lv, A.; Tan, X.; Qin, T.; Liu, T.-Y.; and Yan, R. 2022. Re-creation of creations: A new paradigm for lyric-to-melody generation. *arXiv preprint arXiv:2208.05697*.
- Nichols, E.; Morris, D.; Basu, S.; and Raphael, C. 2009. Relationships between lyrics and melody in popular music. In *ISMIR 2009-Proceedings of the 11th International Society for Music Information Retrieval Conference*, 471–476.
- Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I.; et al. 2018. Improving language understanding by generative pre-training.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140): 1–67.
- Sheng, Z.; Song, K.; Tan, X.; Ren, Y.; Ye, W.; Zhang, S.; and Qin, T. 2021. SongMASS: Automatic Song Writing with Pre-training and Alignment Constraint. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 13798–13805.
- Song, K.; Tan, X.; Qin, T.; Lu, J.; and Liu, T.-Y. 2019. Mass: Masked sequence to sequence pre-training for language generation. *arXiv preprint arXiv:1905.02450*.
- Sun, Y.; Wang, S.; Feng, S.; Ding, S.; Pang, C.; Shang, J.; Liu, J.; Chen, X.; Zhao, Y.; Lu, Y.; et al. 2021. Ernie 3.0: Large-scale knowledge enhanced pre-training for language understanding and generation. *arXiv preprint arXiv:2107.02137*.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, Z.; Liu, H.; Yu, J.; Zhang, T.; Liu, Y.; and Zhang, K. 2024. MuDiT & MuSiT: Alignment with Colloquial Expression in Description-to-Song Generation. *arXiv preprint arXiv:2407.03188*.
- Wang, Z.; and Xia, G. 2021. MuseBERT: Pre-training Music Representation for Music Understanding and Controllable Generation. In *ISMIR*, 722–729.
- Watanabe, K.; Matsubayashi, Y.; Fukayama, S.; Goto, M.; Inui, K.; and Nakano, T. 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)*, 163–172.
- Wu, X.; Huang, Z.; Zhang, K.; Yu, J.; Tan, X.; Zhang, T.; Wang, Z.; and Sun, L. 2023. MelodyGLM: Multi-task Pre-training for Symbolic Melody Generation. *arXiv preprint arXiv:2309.10738*.
- Xiao, D.; Li, Y.-K.; Zhang, H.; Sun, Y.; Tian, H.; Wu, H.; and Wang, H. 2021. Ernie-gram: Pre-training with explicitly n-gram masked language modeling for natural language

- understanding. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1702–1715. Association for Computational Linguistics.
- Yang, L.-C.; and Lerch, A. 2020. On the evaluation of generative models in music. *Neural Computing and Applications*, 32(9): 4773–4784.
- Yu, J.; Wu, S.; Lu, G.; Li, Z.; Zhou, L.; and Zhang, K. 2024. Suno: potential, prospects, and trends. *Frontiers of Information Technology & Electronic Engineering*, 1–6.
- Yu, Y.; Srivastava, A.; and Canales, S. 2021. Conditional lstm-gan for melody generation from lyrics. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 17(1): 1–20.
- Zeng, M.; Tan, X.; Wang, R.; Ju, Z.; Qin, T.; and Liu, T.-Y. 2021. Musicbert: Symbolic music understanding with large-scale pre-training. 791–800.
- Zhang, C.; Chang, L.; Wu, S.; Tan, X.; Qin, T.; Liu, T.-Y.; and Zhang, K. 2022. Relyme: Improving lyric-to-melody generation by incorporating lyric-melody relationships. In *Proceedings of the 30th ACM International Conference on Multimedia*, 1047–1056.
- Zhang, K.; Wu, X.; Zhang, T.; Huang, Z.; Tan, X.; Liang, Q.; Wu, S.; and Sun, L. 2023. WuYun: exploring hierarchical skeleton-guided melody generation using knowledge-enhanced deep learning. *arXiv preprint arXiv:2301.04488*.
- Zheng, E.; Moh, M.; and Moh, T.-S. 2017. Music genre classification: A n-gram based musicological approach. In *2017 IEEE 7th International Advance Computing Conference (IACC)*, 671–677. IEEE.