

Translate the Beauty in Songs: Jointly Learning to Align Melody and Translate Lyrics

Chengxi Li^{1*}, Kai Fan^{2*}, Jiajun Bu¹, Boxing Chen², Zhongqiang Huang², Zhi Yu^{1†}

¹Zhejiang Provincial Key Laboratory of Service Robot, Zhejiang University

²Alibaba DAMO Academy

{chengxili, yuzhirenzhe@zju.edu.cn, bjj@cs.zju.edu.cn}

{k.fan,boxing.cbx,z.huang}@alibaba-inc.com

Abstract

Song translation requires both translation of lyrics and alignment of music notes so that the resulting verse can be sung to the accompanying melody, which is a challenging problem that has attracted some interests in different aspects of the translation process. In this paper, we propose Lyrics-Melody Translation with Adaptive Grouping (LTAG), a holistic solution to automatic song translation by jointly modeling lyrics translation and lyrics-melody alignment. It is a novel encoder-decoder framework that can simultaneously translate the source lyrics and determine the number of aligned notes at each decoding step through an adaptive note grouping module. To address data scarcity, we commissioned a small amount of training data annotated specifically for this task and used large amounts of augmented data through back-translation. Experiments conducted on an English-Chinese song translation data set show the effectiveness of our model in both automatic and human evaluations.¹

1 Introduction

Song translation is a meaningful human endeavor to climb high Tower of Babel for inter-culture exchange. Yet it has not received much attention in the natural language processing (NLP) community despite the advancement of machine translation technologies, especially Neural Machine Translation (NMT) (Bahdanau et al., 2015; Vaswani et al., 2017; Hassan et al., 2018), and the expanding interests of solving real-world problems using artificial intelligence techniques. Challenges include the lack of efficient means to collect parallel lyrics and alignment data, the difficulty of modeling the complex interaction between texts and melody and

* Equal contribution. Work was done during Chengxi Li’s research internship at DAMO Academy, Alibaba Group.

† Corresponding author.

¹The audio and score samples can be found at [LTAG2023.github.io](https://github.com/LTAG2023)

But you play — it to the bea - t

But	you	play	it	to	the	beat					
G4	G4	A4	A4	A4	G4	G4	A4	A4	G4	F4	E4
$\frac{1}{16}$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{1}{16}$	$\frac{1}{2}$
但	你	没	有	好	好	珍					惜

但你没——有 好好珍——惜

Figure 1: An example of the comprehensive translation for “But you play it to the beat” in *Rolling In the Deep*.

imperceptive evaluation of scores. While closely related to text translation, song translation is a more involved task. In addition to the general considerations of word choice and word order in translation, human translators of songs need to have a mastery of cultural traditions and the poetic usage of both source and target languages. Furthermore, the translated lyrics need to be properly aligned with the melody, as shown in Figure 1, to maintain the intact beauty of the song, a factor that is indispensable in song translation (Franzon, 2015).

Researchers have explored Singing Voice Synthesis (SVS) (Liu et al., 2022a,c,b) to automate the vocal singing of songs given the input lyrics and scores, which laid the foundation of convenient and perceptive evaluation and a prospective empirical usage of automatically generated songs. However, there is very few previous studies in the direction of Automatic Song Translation (AST). The sole work (Guo et al., 2022) we are aware of focuses on matching tones and rhythms for the translated target words for tonal languages, by imposing constraints during NMT inference. Their direct use of text translation models and strict mapping between notes and tokens, however, is unable to capture the more involved nature of song translation. While the number of notes provides an easy upper bound on the length of translation, the delicate

alignment between lyrics and melody, as observed in [Haapaniemi and Laakkonen \(2019\)](#), should not be dictated solely by simple rigid rules.

In this paper, we propose Lyrics-Melody Translation with Adaptive Grouping (LTAG), the first comprehensive solution to the AST problem, by jointly modeling of both lyrics translation and lyrics-melody alignment within the transformer-based encoder-decoder framework. LTAG incorporates both lyrics and melody in an end-to-end manner and employs an adaptive grouping module to explicitly model the alignment between lyrics and melody. To facilitate training, we produce the first (Chinese-English) bilingual lyrics-melody alignment data set. To address the data scarcity problem, we also generate a large amount of bilingual lyrics-melody data through back-translation of monolingual lyrics-melody alignment data, which is used together with the high quality manual annotations through a curriculum training strategy. Our experiments show that songs translated by LTAG are both faithful to the original lyrics and singable to the melody, as measured by both automatic metrics and human judges majoring in music. Main contributions of this work are as follows:

- (1) We propose the first joint lyrics translation and lyrics-melody alignment framework LTAG to solve the AST task in a comprehensive manner.
- (2) We design an adaptive grouping method for monotonic lyrics-melody alignment prediction that helps achieve high-quality lyrics translation and to provide flexible and reasonable lyrics-to-melody alignments in the same auto-regressive process.
- (3) We produce the first bilingual lyrics-melody alignment data set that will be released publicly to facilitate further research in this field. We also leverage the back-translation and the curriculum learning strategy to boost performance.
- (4) Our experiments show that LTAG outperforms baselines by a notable margin. Human evaluations indicate that our proposed flexible alignments together with lyrics translation achieves satisfying song translation results.

2 Related Works

Lyrics and Song translations have recently drawn attention from the NLP community. Automatic lyrics translation has been approached by rule-based methods, statistical machine translation methods, finite-state methods with rhythmic and lexical constraints ([Gervás, 2002](#); [Manurung, 2004](#);

[He et al., 2012](#)), and more recently by neural methods ([Ghazvininejad et al., 2016, 2017, 2018](#)). Traditional song translation research has made progress in terms of lyrics translation and lyrics-melody alignment through linguistic knowledge ([Haapaniemi and Laakkonen, 2019](#); [Low, 2003, 2008, 2022](#); [Franzon, 2015](#); [Desblache, 2018](#)). Often, the object of these research was artificial songs. These methods pursue lyrics-melody alignment and lyrics translation in separate tracks. [Guo et al. \(2022\)](#) pursue song translation as a type of Constrained Text Translation. Previous works ([Hokamp and Liu, 2017](#); [Lakew et al., 2019](#); [Li et al., 2020](#); [Zou et al., 2021](#)) on constraining the decoding process are shown to be effective in performance and convenient in implementation. Others that impose constraints during training, such as adding format embedding ([Li et al., 2020](#)), introducing special tags and rescoring length control ([Lakew et al., 2019](#); [Saboo and Baumann, 2019](#)), are data-driven methods and show good performance. In this paper, we propose the lyrics translation model with lyrics-melody alignments for domain shift and length control, and overcome the problem of domain mismatch and data sparsity by using monolingual data.

Lyrics Generation with Alignment Prediction, one of the most important tasks in automatic song production, has received much attention recently. In recent works of machine translation, graph neural networks (GNNs) are widely used for alignment prediction between the source and target sentences ([Li et al., 2022, 2023](#)). In our researched domain, most of the current works ([Lee et al., 2019](#); [Chen and Lerch, 2020](#); [Sheng et al., 2021](#); [Ju et al., 2021](#); [Ma et al., 2021](#); [Xue et al., 2021](#)) adopt the sequence generation method, but with different objectives. Some constrain rhythmic alignment, others theme and target genre. Other works ([Sheng et al., 2021](#); [Ju et al., 2021](#)) apply the attention mechanism and find the lyrics-melody alignment via dynamic programming on the attention weights matrix. This method sometimes results in non-monotonic output. Most importantly, it seems that their alignment component is akin to a post-processing module rather than an integrated unit that learns the dynamic alignment such that the lyrics generation is constrained. In our proposal, we take advantage of the monotonic nature of the lyrics-melody alignments and design a light neural network for alignment prediction in parallel to the translation process.

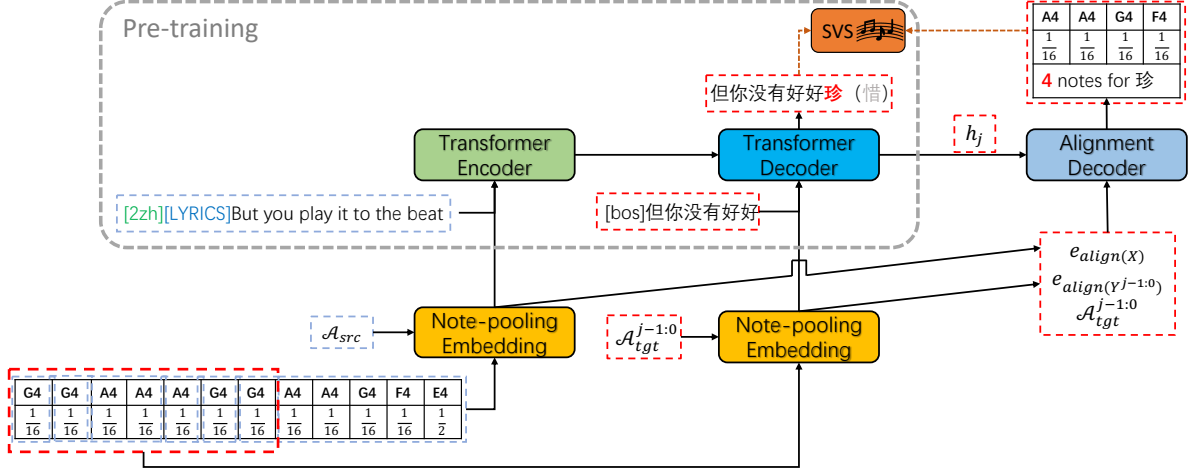


Figure 2: The overall overview of our proposed architecture, illustrated at the j -th decoding step. The transformer decoder will output the target token and the alignment decoder will derive the number of aligned notes.

3 Methodology

In this section, we first describe the LTAG as shown in Figure 2. Then, we detail the adaptive grouping method for alignment prediction and explain how we adapt back-translation for the AST task.

3.1 Overall Architecture

We design an auto-regressive translation architecture that jointly performs lyrics translation and lyrics-melody alignment prediction. As shown in Figure 2, it consists of a transformer-based encoder-decoder pack for lyrics translation, two note-pooling embedding layers that embed and do pooling for notes and alignments, and an alignment decoder. The transformer encoder-decoder is pre-trained with a denoising auto-encoder (Lewis et al., 2020) and for the translation task as in Guo et al. (2022). During pre-training, two prefix tokens indicating the translation direction and the text domain are prepended to the source input. The note-pooling embedding layers shown in Figure 3(a) is a module that processes the melody information. The alignment decoder shown in Figure 3(b) is based on our adaptive notes grouping method that dynamically predicts the number of notes to align to a token during auto-regressive decoding.

3.2 Note-Pooling Embedding

The note-pooling embedding layer takes the notes and alignments as input, and outputs the pooled note embedding and alignment embedding. The input note sequence consists of MIDI pitch and duration of each note. The MIDI pitch and duration can be represented as embedding e_{midi} and e_{dur}

respectively. We define the i -th note embedding:

$$\mathbf{e}_{note}^i = \mathbf{e}_{midi}^i + \mathbf{e}_{dur}^i + \mathbf{e}_p^i \quad (1)$$

where \mathbf{e}_p^i is the positional embedding.

We apply non-overlapping mean-pooling on the note embedding sequence according to the alignment information. Specifically, the embeddings of the consecutive notes that align to the same token are averaged. Mathematically, the alignment information \mathcal{A} is represented as a binary matrix $\mathbf{M} \in \{0, 1\}^{L \times N}$, where L and N denote the sequence length of the tokens and notes. $\mathbf{M}_{ji} = 1$ if the i -th note is aligned to the j -th token. We use \mathbf{M} to efficiently calculate the non-overlapping mean-pooling via matrix multiplication, denoted the result as melody embedding \mathbf{e}_{md} .

$$\mathbf{e}_{md} = \text{Non-Overlap-Mean-Pool}(\mathbf{e}_{note}, \mathbf{M}) \quad (2)$$

The kernel size of this operation is not fixed but varies according to the row sum of \mathbf{M} . The detailed calculation can refer to Appendix A.

Because lyrics-melody alignments are monotonic, we encode the alignment more succinctly by calculating the cumulative sum of the number for aligned notes:

$$\mathbf{s} = \text{CumSum}(\text{RowSum}(\mathbf{M})) \quad (3)$$

where \mathbf{s} is a vector of length L . s^j/N then represents the alignment ratio for each aligned note. We next quantize the cumulative alignment ratios by grouping them into equal-size bins over the range $(0, 1]$, and introduce a set of embedding vectors

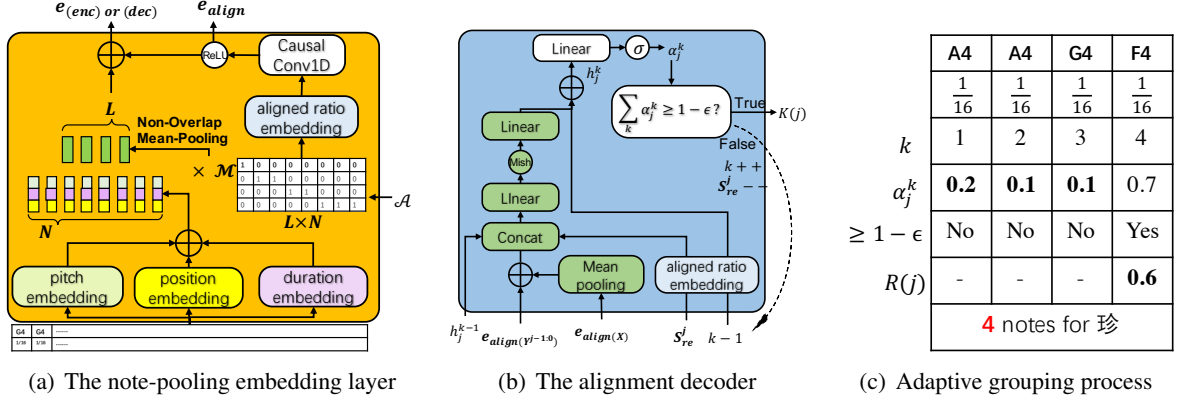


Figure 3: (a) The note-pooling embedding encodes both the note sequence and the alignment information. (b)(c) The alignment decoder computes the number of aligned notes from halting distribution.

\mathbf{E}_{ratio} to represent each bin. Finally, the alignment embedding is calculated as follows.

$$\mathbf{e}_{align}^j = f(\mathbf{E}_{ratio}(s^j/N)) \quad (4)$$

where $f(\cdot)$ is a simple non-linear layer of causal 1D convolution with ReLU activation, and the number of bins is a hyper-parameter. The motivation is to implicitly constrain the translation by the number of aligned notes.

The melody embedding and alignment embedding are summed and then added to the original transformer encoder or decoder input.

$$\mathbf{e}_{enc(dec)} = \mathbf{e}_{token} + \mathbf{e}_p + (\mathbf{e}_{md} + \mathbf{e}_{align}) \quad (5)$$

As calculated in Eq. (2), each melody embedding corresponds to one token. In addition, the causal convolution implies that the alignment embedding tensors also have the same length as the text tokens and guarantees each alignment embedding only observes previous ratio embeddings in an auto-regressive manner. It means that on the decoder, this layer can fit perfectly in the teacher-forcing training.

3.3 Alignment Decoder

Inspired by the Adaptive Computation Time (ACT) (Graves, 2016), we propose the **adaptive grouping** module to model lyrics-melody alignment. As shown in Figure 3(b) and 3(c), this module predicts how many consecutive notes should be assigned to the current token.

For $1 \leq j \leq L_Y$, let y_j be the j -th target token and \mathbf{h}_j be the corresponding hidden state of the last transformer decoder layer. Suppose previous tokens $y_{j-1:0}$ have been aligned to the first $n-1$

notes, we define the following adaptive grouping process by iterating over index k (starting from 1) to derive the number of notes aligned to y_j .

$$\begin{aligned} \mathcal{S}_{re}^j &= N - s_{tgt}^{j-1} \\ \mathbf{h}_j^0 &= \mathbf{h}_j \\ \mathbf{h}_j^k &= g(\mathbf{h}_j^{k-1}, \mathbf{e}_{align(X)}, \mathbf{h}_{align(y_{j-1:0})}, \mathcal{S}_{re}^j, k-1) \\ \alpha_j^k &= \sigma(\text{Linear}(\mathbf{h}_j^k)) \end{aligned}$$

where $\mathbf{e}_{align(X)}$ and $\mathbf{e}_{align(y_{j-1:0})}$ are the alignment embeddings of the full source input and the partial target input respectively, and s_{tgt}^{j-1} is j -th element of vector \mathbf{s} in Eq. (3).

We first calculate the residual number of unaligned notes at the current decoding step j as \mathcal{S}_{re}^j . $\mathbf{e}_{align(X)}$ is fed into an average pooling layer to obtain a single vector, making it always possible to be additive with $\mathbf{e}_{align(y_{j-1:0})}$ of variable length. For all the inputs, we apply a multi-layer network $g(\cdot)$ shown in green in Figure 3(b). Eventually, the sigmoid function $\sigma(\cdot)$ outputs the halting probability α_j^k of the intermediate step. The summation of these probabilities represent the likelihood that the current k notes are aligned to the target token y_j .

Given a hyper-parameter ϵ as a small float number (e.g., 0.01), if $\sum_k \alpha_j^k < 1 - \epsilon$, the adaptive grouping process will continue and re-calculate by incrementing k and decrementing \mathcal{S}_{re}^j . Otherwise, the aligning process halts, and the alignment decoder outputs the number of aligned notes $K(j)$.

$$K(j) = \underset{K}{\operatorname{argmin}} \left\{ \sum_{k=1}^K \alpha_j^k \geq 1 - \epsilon \right\} \quad (6)$$

A positive $\epsilon > 0$ guarantees that $K(j) \geq 1$, i.e., at least one note is aligned. To define the halting

probabilities of $K(j)$ aligned notes, we introduce the remainder $R(j) = 1 - \sum_{k=1}^{K(j)-1} \alpha_j^k$. In this way, α_j^k and $R(j)$ can be valid probability distributions. Figure 3(c) is an example of how the adaptive grouping works.

In the labeled alignment data, the ground truth of the number of aligned notes for each target token is available, denoted as Δ_j . Instead of minimizing the ponder cost $\sum_j K(j) + R(j)$ as in ACT (Graves, 2016), we optimize the following adaptive grouping loss L_G , which could naturally upper bound the token-wise ponder cost via Δ_j .

$$L_G = \left| \sum_j K(j) - N \right| + \sum_j |K(j) - \Delta_j| \\ \approx \left| \sum_j (K(j) - (1 - R(j))) - N \right| \\ + \sum_j |K(j) - (1 - R(j)) - \Delta_j|$$

The variable $K(j)$ is discontinuous with respect to the halting probabilities, so we use $1 - R(j)$ in the approximation to make the loss differentiable (more analysis in Appendix B). Additionally, because tokens aligned to more than one notes are infrequent, we add upweighting to the alignment loss of such tokens for model calibration.

$$L_G = \left| \sum_j (K(j) - 1 + R(j)) - N \right| \\ + \sum_j (|K(j) - 1 + R(j) - \Delta_j| \cdot w_j)$$

where $w_j = 1$ if $\Delta_j = 1$ and $w_j > 1$ is a hyper-parameter if $\Delta_j > 1$.

3.4 Back Translation with Alignments

Although a data set of a few thousand verses with human translation and annotated with alignment information is useful, its quantity is limited. We therefore adopt the widely used back-translation method (Sennrich et al., 2016) to generate more training data. We crawl the web for more available monolingual song data with alignments and build another pre-trained lyrics translation model with length control that is used to back translate the monolingual data into the source language. The length control ensures that the number of tokens is the same as the number of notes after which a one-to-one source-side alignment can be generated. This way, we obtain a comparatively larger data set with noise on the source side but still accurate information on the target side.

Because the back-translated data are much larger than the human annotated one, we in practice design our data loader by following a curriculum

learning way. Initially, the augmented data from back-translation will be mixed with up-sampled the real data from human annotation. In each training epoch, we gradually down-sample the augmented data to raise the ratio of annotated data in the batch. A visualization of the data sampling scheduler is in Figure 6 (See Appendix C).

3.5 Training and Inference

After the pre-training stage, we will optimize the whole model by jointly minimizing the loss from the task of lyrics translation and the task of alignment prediction. Note that the SVS model is pre-trained and only used for evaluation. The overall loss is thus:

$$L_{joint} = \sum_{j=0}^{L_Y} \log P(y_j | y_{j-1:0}, X) + \beta \cdot L_G$$

where β is a hyper-parameter to balance the importance between the two tasks. The inference follows the standard beam search for auto-regressive decoding, while only the last generated token and its corresponding notes should be specially taken care of. Details can be found in Appendix D.

4 Experiments

In this section, we describe the experiment setup, results and analysis on Chinese \leftrightarrow English song translation.

4.1 Experimental Settings

Data Sets

Since there is no publicly available data set with high quality parallel lyrics translation and lyrics-melody alignments, we collect and annotate a data set PopCV (Pop songs with Cover Version) containing both Chinese songs with their English cover version and English ones with Chinese cover version. Since there are no industry standards or published precedence in annotating such a data set, we design an annotation procedure which is time-saving and easy for annotators to carry out. First, we collect the score sheet files of songs from score websites². Then the annotators add lyrics to notes according to how songs are sung in the original and its cover version as conventions³ suggest. We then export the annotated music score files in .musicxml format

²<https://www.musescore.com> and <https://www.midishow.com>

³<https://lilypond.org> and <https://musescore.org/howto>

Models	MOS-T		MOS-S		MOS-Q	
	En→Zh	Zh→En	En→Zh	Zh→En [†]	En→Zh	Zh→En [†]
Human Ref.	3.83 ± 0.06	4.11 ± 0.05	3.92 ± 0.07	\	3.90 ± 0.06	\
GagaST	3.66 ± 0.06	3.72 ± 0.05	3.49 ± 0.10		3.65 ± 0.05	
LTAG-cl	3.66 ± 0.05	3.79 ± 0.05	3.58 ± 0.07		3.62 ± 0.05	
only bt	3.69 ± 0.05	3.80 ± 0.04	3.53 ± 0.09		3.63 ± 0.05	
w/o bt	3.64 ± 0.05	3.30 ± 0.05	2.16 ± 0.05		3.14 ± 0.04	
LTAG	3.71 ± 0.05	3.85 ± 0.05	3.68 ± 0.05		3.69 ± 0.04	
only bt	3.71 ± 0.05	3.80 ± 0.05	3.58 ± 0.07		3.65 ± 0.04	
w/o bt	3.69 ± 0.05	3.28 ± 0.04	3.63 ± 0.07	3.67 ± 0.04		

Table 1: The Mean Opinion Score in translation intelligibility and naturalness (MOS-T), singability (MOS-S) and overall quality (MOS-Q) with 95% confidence intervals. The translation direction with [†] means that audio samples of the translated song for evaluation are generated with the voice synthesis model that is not trained for that target language. So those results are presented in Appendix F and for reference only.

	Lang	Songs	Lyrics	Source
LMD	En	\	152,991	Back Translation
Changba	Zh	\	542,034	Back Translation
PopCV	En,Zh	79	2,959	Annotation
testset	En,Zh	25	629	Annotation

Table 2: Statistics of datasets in our experiments

and automatically extract lyrics and their aligned notes. Please refer to Appendix E for details.

For the data used in back-translation, we use LMD⁴ (Yu et al., 2021) for English songs with alignments to melody and a data set crawled from Changba App for Chinese songs. We first pre-train two lyrics translation models with length control, one in each direction, and then translate the above two data sets. The translated lyrics are one-on-one aligned to the notes. Two sets of back-translated data are used for training only while testing is done on real data with human annotations. An overview of the data is in Table 2.

Evaluation Metrics

The most convincing evaluation of how our model works is whether the translated songs can be sung, understood, and, most importantly, enjoyed. Thus, we follow Sheng et al. (2021) and show annotators the resulting score of the song with translated lyrics. To verify the singability in the end-to-end manner, we additionally use an open-source Chinese singing voice synthesis (SVS) model (Liu et al., 2022a) to supply the annotators with an actual audio rendition of the songs for more intuitive feeling.

⁴<https://github.com/yyllab/Lyrics-Conditioned-Neural-Melody-Generation>

We randomly select 20 verses from the test set and show the music sheets and synthesized singing voice (see Appendix E) of each translated verse to five annotators. For automatic evaluations, we use sacreBLEU⁵. For translation intelligibility, naturalness, singability and overall quality evaluation, we use mean opinion score (MOS) in human evaluations, referred to MOS-T, MOS-S and MOS-Q. In evaluating the alignments, the traditional AER does not apply here because in addition to machine-produced alignments, the target translation is also machine-produced. Instead, we propose an Alignment Score (AS) that calculates the weighted intersection over ground truth (IOG) of the empirical probability density between the predicted and the true alignments:

$$AS = \frac{\sum_k \min(\text{freq}_{pred}^k / F_{pred}, \text{freq}_{gt}^k / F_{gt}) * k}{\sum_k (\text{freq}_{gt}^k / F_{gt}) * k} \quad (7)$$

where k represents the number of aligned notes, and $F = \sum_k \text{freq}^k$. More details are included in Appendix F.

Model Configurations

The token embeddings of the Transformer encoder and decoder have dimension of 256 and are shared. In the note-pooling embedding layer, the size of the lookup table for MIDI pitch and duration type are set to 128 and 31. The halting hyper-parameter epsilon ϵ for the adaptive grouping process is 0.05. w_j is 5 when $\Delta_j > 1$, and β is 0.8 in the joint loss L_{joint} . The beam size during decoding is 5.

The LTAG model is pre-trained on the WMT data and the crawled lyrics data, including the par-

⁵<https://github.com/mjpost/sacrebleu>




Figure 4: The overlapped histograms of ground truth alignments and predicted alignments on En→Zh test set.

Models	BLEU↑		AS. ↑	
	En→Zh	Zh→En	En→Zh	Zh→En
GagaST	11.87	5.67	0.701	0.468
LTAG-cls	14.21	10.01	0.827	0.555
only bt	15.54	10.21	0.709	0.667
w/o bt	13.73	8.26	0.704	0.490
LTAG	16.02*	10.68	0.923	0.781
only bt	16.27	10.26*	0.880*	0.718*
w/o bt	14.12	7.86	0.845	0.710
w/o e_{align}	15.16	9.24	0.852	0.703

Table 3: The sacreBLEU and Alignment Score on both translation directions. * means the second highest result within the row.

allel and the monolingual corpora. The sampling ratio scheduling of augmented data and annotated data are described in Appendix C.

For voice synthesis, we convert the Chinese lyrics into phonemes by *pinyin* (Ren et al., 2020) and set the hop size and frame size to 128 and 512 for the sample rate of 24kHz. Pitch inputs to the SVS model are all re-tuned to the range between *A3* and *C5* in C major. Besides, we apply some post-processing in inference to generate scores and singing voice for more tolerance (Appendix D).

4.2 Main Results

We compare LTAG with two baseline systems. One is the GagaST system (Guo et al., 2022), which focuses on the tonal aspect of Chinese. The other one is a variation of our model. This variation uses a transformer-layer based classifier (LTAG-cls) instead of our alignment decoder to predict the number of aligned notes. The maximum number of aligned notes is 30, the same as allowed maximum $K(j)$ in alignment decoder. Besides, we show results from the human reference.

4.2.1 Translation Evaluation

We first report the human evaluation metrics (MOS-T) on both Chinese-to-English (Zh→En) and English-to-Chinese (En→Zh) song translation tasks in Table 1. LTAG generally gains improvements among all systems while the gap between different systems and settings is not obvious. It’s partly because the lyrics translation by professionals is usually free translation rather than literal translation. A missing word in different slices can cause negative, neutral or even positive effect. Only obvious semantic deviations or grammatical mistakes lead to certain score decrease. As discussed in MOS-T, automatic metric BLEU may not be a good criterion to compare the machine translation and free translation for lyrics. But we still present the BLEU results in Table 3. We can see that our proposed system LTAG significantly outperforms the recent baseline GagaST by a large margin on both translation directions. As to the variant model LTAG-cls, the LTAG is still slightly better.

4.2.2 Lyrics-Melody Alignment Evaluation

As for lyrics-melody alignment quality, we report the human evaluation metrics (MOS-S) on en-zh translation direction. In Table 1, LTAG considerably outperforms other systems, especially better than the GagaST with simple length control decoding. Notably, the variant version LTAG cls performs worse than other systems, which indicates that more flexible alignments between lyrics and melody bring listening enjoyment to audience when it’s reasonable enough. Otherwise, the flexibility may be counterproductive. We also evaluate the alignment quality by using the histograms of the number of aligned notes in Figure 4. In Table 3, we calculate the Alignment Score between the histograms of each system and the true histograms. The histograms show that the distribution of align-

(a) Source and Reference. Left: En→Zh. Right: Zh→En

(b) GagaST

(c) LTAG-cls

(d) LTAG

Figure 5: Example scores of the source, reference and the translation for “Is love I can be sure of” in *Will You Love Me Tomorrow* and “tā hui yǒu duō xìng yùn” in *Xiǎo Xīng Yùn* from three systems.

ments generated by LTAG resemble those of the true alignments while “GagaST” lacks variety by providing only one-on-one alignments between the lyrics and melody. In conclusion, both results demonstrate that the adaptive grouping method shows significant advantage over the length control or simple classifier in predicting reasonable alignment between the translated lyrics and melody.

In Table 1, MOS-Q mainly reflects the overall intelligibility, naturalness, singability and beauty of the song translation. Since the translation and alignment quality both contribute to the final result, the difference between methods seem less visible. But considering the 95% confidence, we can conclude that the LTAG still ranks best.

4.3 Ablation Study and Analysis

We first conduct ablation experiments to study the effects of back translation data with various settings. In Table 1, we have the following findings for LTAG and LTAG-cls. (1) Since the back-translation data is obviously larger than the real annotation data, there is almost no difference if only back-translation data is used for training. This enables the possibility of training our model in unsupervised way. (2) If only the limited supervised data is

used, the performance apparently becomes worse. (3) LTAG is consistently better than LTAG-cls in all ablation experiments. In addition, we verify the importance of the novel alignment embedding e_{align} by removing it from the note-pooling embedding layer and alignment decoder, and observe a non-negligible decrease on both BLEU and AS.

Some case studies in Figure 5 suggest that, when the tokens in lyrics fall into one-to-many alignments, GagaST usually provides inappropriate lyrics translation or even decodes non-vocal tokens such as comma to meet the length constraint. It will hurt both the translation quality and the singability of the translated lyrics. In contrast, the simple classifier following transformer layers is enough for flexible alignments. However, our evaluation results indicate our light weighted alignment decoder is capable of providing delicate alignments between tokens and notes.

5 Conclusion

In this work, we propose LTAG, a lyrics translation model with lyrics-melody alignments that allows simultaneous generation of target text and alignment to the music notes. We propose an adaptive grouping method that fits in the auto-regressive

translation process. To better train and evaluate our model, we also annotate a new song translation data set PopCV containing English and Chinese songs with their cover version in both languages and with lyrics-melody alignments. For training, we also employ back-translation to leverage the more abundantly available monolingual lyrics data with lyrics-melody alignments in a curriculum learning way. Evaluations with both the automatic and human metrics show that LTAG is capable of producing natural, singable and enjoyable translation results.

Concerns of the Ethical Impacts

This work develops a possible automatic method for song translation. Therefore, if we release our repository and data set, there is the potential of abusing to synthesize score sheets and texts, especially may cause copyright issues. Thus, we choose the dataset license: CC by-nc-sa 4.0. In this paper, we thoroughly discuss strengths and shortcomings of our proposed model and perform a series of experiments to support them. Codes, model checkpoints and data set will be released upon acceptance after desensitization and compliance examination.

Limitations and Future Work

Our work mainly focuses on Chinese and English songs while languages from other language families are not involved. They may lack lyrics and song data to perform data augmentations or to train a SVS system used in our human-evaluation process. Future endeavors may lie in mining and utilizing song translation data from richer sources and more languages.

Acknowledgments

We would like to thank all the anonymous reviewers for the insightful and helpful comments.

This work is supported by the National Natural Science Foundation of China (Grant No. 61972349), the National Key R&D Program of China (No. 2018YFC2002603), Alibaba-Zhejiang University Joint Institute of Frontier Technologies, and Alibaba Group through Alibaba Research Intern Program.

References

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of the*

Inter-national Conference on Learning Representations.

Yihao Chen and Alexander Lerch. 2020. Melody-conditioned lyrics generation with seqgans. *2020 IEEE International Symposium on Multimedia (ISM)*, pages 189–196.

Lucile Desblache. 2018. Translation of music. *An encyclopedia of practical translation and interpreting*, pages 297–324.

Johan Franzon. 2015. Three dimensions of singability. an approach to subtitled and sung translations. *Text and Tune. On the Association of Music and Lyrics in Sung Verse*. Bern: Peter Lang, pages 333–346.

Pablo Gervás. 2002. Wasp: Evaluation of different strategies for the automatic generation of spanish verse.

Marjan Ghazvininejad, Yejin Choi, and Kevin Knight. 2018. [Neural poetry translation](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 67–71, New Orleans, Louisiana. Association for Computational Linguistics.

Marjan Ghazvininejad, Xing Shi, Yejin Choi, and Kevin Knight. 2016. [Generating topical poetry](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1183–1191, Austin, Texas. Association for Computational Linguistics.

Marjan Ghazvininejad, Xing Shi, Jay Priyadarshi, and Kevin Knight. 2017. [Hafez: an interactive poetry generation system](#). In *Proceedings of ACL 2017, System Demonstrations*, pages 43–48, Vancouver, Canada. Association for Computational Linguistics.

Alex Graves. 2016. Adaptive computation time for recurrent neural networks. *arXiv preprint arXiv:1603.08983*.

Fenfei Guo, Chen Zhang, Zhirui Zhang, Qixin He, Kejun Zhang, Jun Xie, and Jordan Boyd-Graber. 2022. [Automatic song translation for tonal languages](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 729–743, Dublin, Ireland. Association for Computational Linguistics.

Riku Haapaniemi and Emma Laakkonen. 2019. The materiality of music: interplay of lyrics and melody in song translation. *Translation Matters*.

Hany Hassan, Anthony Aue, Chang Chen, Vishal Chowdhary, Jonathan Clark, Christian Federmann, Xuedong Huang, Marcin Junczys-Dowmunt, William Lewis, Mu Li, et al. 2018. Achieving human parity on automatic chinese to english news translation. *arXiv preprint arXiv:1803.05567*.

- Jing He, Ming Zhou, and Long Jiang. 2012. [Generating chinese classical poems with statistical machine translation models](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 26(1):1650–1656.
- Chris Hokamp and Qun Liu. 2017. [Lexically constrained decoding for sequence generation using grid beam search](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1535–1546, Vancouver, Canada. Association for Computational Linguistics.
- Zeqian Ju, Peiling Lu, Xu Tan, Rui Wang, Chen Zhang, Songruoyao Wu, Kejun Zhang, Xiangyang Li, Tao Qin, and Tie-Yan Liu. 2021. [Telemelody: Lyric-to-melody generation with a template-based two-stage method](#). *CoRR*, abs/2109.09617.
- Surafel Melaku Lakew, Mattia Di Gangi, and Marcello Federico. 2019. [Controlling the output length of neural machine translation](#). In *Proceedings of the 16th International Conference on Spoken Language Translation*, Hong Kong. Association for Computational Linguistics.
- Hsin-Pei Lee, Jih-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.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Lei Li, Kai Fan, Hongjia Li, and Chun Yuan. 2022. [Structural supervision for word alignment and machine translation](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 4084–4094.
- Lei Li, Kai Fan, Lingyu Yang, Hongjia Li, and Chun Yuan. 2023. [Neural machine translation with dynamic graph convolutional decoder](#). *arXiv preprint arXiv:2305.17698*.
- 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.
- Jinglin Liu, Chengxi Li, Yi Ren, Feiyang Chen, and Zhou Zhao. 2022a. [Diffsinger: Singing voice synthesis via shallow diffusion mechanism](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11020–11028.
- Jinglin Liu, Chengxi Li, Yi Ren, Zhiying Zhu, and Zhou Zhao. 2022b. [Learning the beauty in songs: Neural singing voice beautifier](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7970–7983, Dublin, Ireland. Association for Computational Linguistics.
- Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. 2022c. [Pseudo numerical methods for diffusion models on manifolds](#). In *International Conference on Learning Representations*.
- Peter Low. 2003. [Singable translations of songs](#). *Perspectives: Studies in Translatology*, 11:87–103.
- Peter Low. 2008. [Translating songs that rhyme](#). *Perspectives: Studies in translatology*, 16(1-2):1–20.
- Peter Low. 2022. [Translating the Texts of Songs and Other Vocal Music](#), Cambridge Handbooks in Language and Linguistics, page 499–518. Cambridge University Press.
- Xichu Ma, Ye Wang, Min-Yen Kan, and Wee Sun Lee. 2021. [Ai-lyricist: Generating music and vocabulary constrained lyrics](#). In *Proceedings of the 29th ACM International Conference on Multimedia*, MM ’21, page 1002–1011, New York, NY, USA. Association for Computing Machinery.
- Hisar Maruli Manurung. 2004. [An evolutionary algorithm approach to poetry generation](#).
- Yi Ren, Xu Tan, Tao Qin, Jian Luan, Zhou Zhao, and Tie-Yan Liu. 2020. [Deepsinger: Singing voice synthesis with data mined from the web](#). In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1979–1989.
- Ashutosh Saboo and Timo Baumann. 2019. [Integration of dubbing constraints into machine translation](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 94–101, Florence, Italy. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. 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](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13798–13805.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all](#)

you need. *Advances in neural information processing systems*, 30.

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, Abhishek Srivastava, and Simon Canales. 2021. [Conditional lstm-gan for melody generation from lyrics](#). *ACM Trans. Multimedia Comput. Commun. Appl.*, 17(1).

Xu Zou, Da Yin, Qingyang Zhong, Hongxia Yang, Zhilin Yang, and Jie Tang. 2021. [Controllable generation from pre-trained language models via inverse prompting](#). In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery; Data Mining, KDD '21*, page 2450–2460, New York, NY, USA. Association for Computing Machinery.

A Pooling Matrix in the Note-pooling Embedding Layer

We have note embedding $\mathbf{e}_{note} \in \mathbb{R}^{N \times d}$ (d is the embedding dimension) and alignment matrix $\mathbf{M} \in \{0, 1\}^{L \times N}$. The non-overlapped mean-pooling can be calculated as follows.

$$\mathbf{W} = \mathbf{M} / \text{sum}(\mathbf{M}, \text{dim} = -1, \text{keepdim} = \text{True})$$

$$\mathbf{e}_{md} = \mathbf{W} * \mathbf{e}_{note}$$

where $/$ is element-wise division and $*$ is matrix multiplication. By leveraging `gather` and `scatter` operations, the non-overlapped mean-pooling can even be computed in batch.

B Analysis of Adaptive Grouping Loss

By the definition of the adaptive grouping loss, we only need to analyze the following term.

$$|K(j) - (1 - R(j)) - \Delta_j|$$

If $K(j) > \Delta_j$ in the forward pass, we have $K(j) - \Delta_j \geq 1$ because they are both positive integers. In order to encourage the loss to become smaller, $1 - R(j) = \sum_{k=1}^{K(j)-1} \alpha_j^k$ should become larger. In other words, the optimization will push $\sum_{k=1}^{K(j)-1} \alpha_j^k$ to be larger towards $K(j) - \Delta_j$. Note that the theoretical upper bound of $\sum_{k=1}^{K(j)-1} \alpha_j^k$ is $K(j) - 1$, which is larger or equal to $K(j) - \Delta_j$. Thus, this optimization is possible and it will meet the following condition during optimization.

$$\sum_{k=1}^{K(j)-1} \alpha_j^k \geq 1 - \epsilon.$$

By definition of $K(j)$, we have the following conclusion.

$$K(j)^{\text{new}} = \arg \min_K \left\{ \sum_{k=1}^K \alpha_j^k \geq 1 - \epsilon \right\} \leq K(j) - 1$$

If $K(j) < \Delta_j$, a similar analysis can be derived. $\sum_{k=1}^{K(j)-1} \alpha_j^k \rightarrow 0$ should be encouraged to pursue a smaller loss. It implies if the $K(j)$ -th halting probability doesn't satisfy the condition $\alpha_j^{K(j)} \geq 1 - \epsilon$, the $K(j)^{\text{new}}$ will have an increasing trend. However, if $\alpha_j^{K(j)} \geq 1 - \epsilon$, the optimization will be stuck. We may optimize $|K(j) - (1 - R(j) + \alpha_j^{K(j)}) - \Delta_j|$, *i.e.*, $|K(j) - \sum_{k=1}^{K(j)} \alpha_j^k - \Delta_j|$. In practice, we found

this is a rare case and it will completely disappear after several epochs. So we adopt the unified adaptive grouping loss.

If $K(j) = \Delta_j$, it means we can safely remove this term in the loss.

C Scheduler of Curriculum Learning

The down sampling ratio of back translation data starts at 1.00 and decrease to 0.01 at the half of total training epochs. The sampling ratio of annotation data starts at 20.00 for upsampling and decrease to 5.00 at the end of total training epochs.




Figure 6: An illustration of how we use back translation data together with annotated data in co-translation training. “bt” represent data from back translation data augmentation and “at” represent data from annotation.

D Post-processing In Inference

In order to generate scores and singing voice in line with musical rules, we add some rule-based post-processing to the alignment predictions for more tolerance. For cases where total number of aligned notes is larger than the number of notes in the melody, we simply truncate the predicted number of aligned notes from the last token to the first or from the first token to the last. For cases of fewer number of predicted notes, we add the number of difference all to the last token.

E Data Annotation and Human Evaluation

Annotators are students who major in music, vocal singing or relevant specialty. They all speak bilingual languages with Chinese and English, so they are also qualified for translation quality evaluation. For data annotation and human evaluation, each person gets reasonably paid according to the

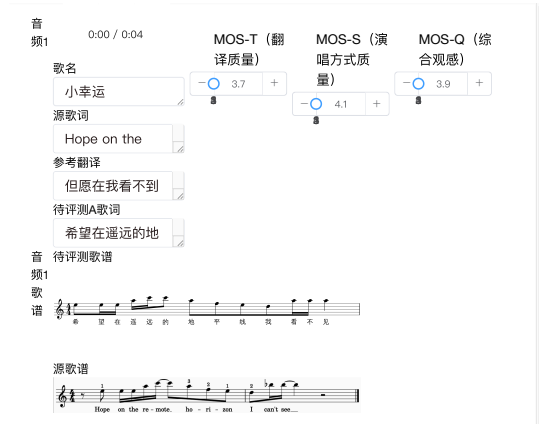


Figure 7: An example of evaluation front-end interface for human evaluation.

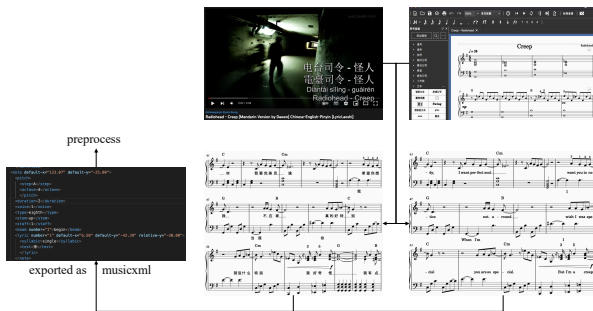


Figure 8: An overview of our propdata annotated notation pipeline.

individual workload. The annotation guidance and evaluation guidance can be found in supplement materials. Figure 7 is an example of visual front-end interface for human evaluation. The pipeline for data annotation is shown in Figure 8

F Song Translation Evaluations

Firstly, we clarify the signature⁶ for sacreBLEU in our evaluations. Notably, we take the lyrics that are both human-translated and can be sung as the single golden reference in our sacreBLEU score evaluations. Other human-translated lyrics are not guaranteed to be able to be sung. So the comparison of sacreBLEU scores including them may be unfair.

Besides, here we show the MOS-S and MOS-Q for Chinese-to-English song translation for reference. Lack of open data set for English singing voice synthesis caused the bad quality of synthesized English singing voice in inference. So we have to use the Chinese SVS system to synthesize the English songs. According to the feedback from

annotators, this gap influence their feeling about translation results to some extent. As is shown in Table 4, results significantly drop compared to those of En→Zh. So we leave this to appendix part for reference.

	MOS-S	MOS-Q
	Zh→En	Zh→En
Human Ref.	4.05 ± 0.08	4.06 ± 0.07
GagaST	3.04 ± 0.02	3.08 ± 0.03
LTAG-clc	3.03 ± 0.02	3.10 ± 0.01
only bt	3.11 ± 0.04	3.15 ± 0.03
w/o bt	3.27 ± 0.04	3.20 ± 0.02
LTAG	3.39 ± 0.03	3.61 ± 0.05
only bt	3.39 ± 0.05	3.60 ± 0.02
w/o bt	3.29 ± 0.05	3.32 ± 0.02

Table 4: The Mean Opinion Score singability (MOS-S) and overall quality (MOS-Q) for Zh→En samples with 95% confidence intervals.

⁶nrefs:1lcase:mixedleff:noltok:13alsmooth:explversion:2.0.0