

A COMPUTATIONAL EVALUATION FRAMEWORK FOR SINGABLE LYRIC TRANSLATION

Haven Kim¹ Kento Watanabe² Masataka Goto² Juhan Nam¹

¹ Graduate School of Culture Technology, KAIST, South Korea

² National Institute of Advanced Industrial Science and Technology (AIST), Japan

khaven@kaist.ac.kr, kento.watanabe@aist.go.jp, m.goto@aist.go.jp, juhan.nam@kaist.ac.kr

ABSTRACT

Lyric translation plays a pivotal role in amplifying the global resonance of music, bridging cultural divides, and fostering universal connections. Translating lyrics, unlike conventional translation tasks, requires a delicate balance between singability and semantics. In this paper, we present a computational framework for the quantitative evaluation of singable lyric translation, which seamlessly integrates musical, linguistic, and cultural dimensions of lyrics. Our comprehensive framework consists of four metrics that measure syllable count distance, phoneme repetition similarity, musical structure distance, and semantic similarity. To substantiate the efficacy of our framework, we collected a singable lyrics dataset, which precisely aligns English, Japanese, and Korean lyrics on a line-by-line and section-by-section basis, and conducted a comparative analysis between singable and non-singable lyrics. Our multidisciplinary approach provides insights into the key components that underlie the art of lyric translation and establishes a solid groundwork for the future of computational lyric translation assessment.

1. INTRODUCTION

Translating lyrics is a prevailing method of enhancing the global appeal and allure of music across a multitude of genres, such as theatre music, animation music, pop music, etc [1]. Furthermore, recent advancements in media technology have facilitated the exchange of intercultural products and globalized fandom culture, resulting in an increase in the popularity of user-translated lyrics across diverse social media platforms [2].

Despite its popularity, lyric translation is acknowledged as a challenging field, requiring an interdisciplinary approach [2]. As early as 1915, it was suggested that an ideal lyric translator should possess expertise in both linguistics and music, highlighting the need for a comprehensive understanding of the principles and techniques used in translation studies, coupled with a background in musicology [3].

Moreover, it is crucial to understand the cultural context of each language, such as different strategies employed for forming rhymes [4]. Because of these challenges, the systematic analysis and evaluation of lyric translation remain an under-researched topic of study. Thus far, only a few have proposed guidelines for scoring the quality of translated lyrics [4, 5]. While these principled approaches have proven successful, they lack automation. Consequently, despite the growing interest in the development of neural lyric translation, its evaluation has predominantly relied on human evaluation, making the evaluation process time-consuming, unreliable, and subjective [6–8].

Our study aims to computationally analyze and evaluate lyric translation based on a comprehensive understanding of lyrics that accounts for their musical, linguistic, and cultural elements. Unlike prior research that only proposed rhyme-scoring guidelines applicable to English [4], our framework is extendable to Japanese and Korean. Though our framework may be limited in its application to specific languages, we strive to provide valuable insights into establishing distinct evaluation rules for phoneme repetition in diverse languages. Our comprehensive framework employs a multifaceted evaluation approach that examines lyric translation from four distinct perspectives: syllable counts, phoneme repetition, musical structure, and semantics. In the remainder of this paper, we explicate the rationale behind our selection of these perspectives by delving into the unique characteristics of lyric translation that differentiate it from general language translation tasks. In addition, we introduce the singable lyrics dataset we collected, which features line-by-line and section-by-section alignment of English, Japanese, and Korean lyrics. Moving forward, we propose robust evaluation metrics for lyric translation and analyze the results of our experiments based on the perspectives mentioned above. Finally, we conclude our paper by reflecting on the profound insights gleaned from our experiment and highlighting possible directions for future research.

2. BACKGROUND

Previous research indicates that linguistic analysis methods designed for standard text may not achieve desired outcomes when used to examine lyrics [9]. Although automated evaluation metrics, such as n -gram-based [10–12] or neural approaches [13], have proven valuable and effective in assessing conventional machine translation tasks, they



fall short in evaluating lyric translation. This is due to the unique characteristics of lyrics that render the translation process subject to many constraints and less direct [14].

One of the most significant constraints is the syllable count. This is because the original and translated lyrics must match the same melody lines, while it is a common practice to tweak the melody to accommodate minor changes in syllable count [4, 15]. In fact, conveying the same message in different languages requires vastly different syllable counts. For example, “Happy New Year” in English consists of 4 syllables, whereas 15 and 9 are required for Japanese (あけましておめでとうございます) and Korean (새해 복 많이 받으세요), respectively. For the numerical comparison, we examined PAWS-X, a dataset that contains 23,659 English sentences paired with human-translated sentences in various languages [16]. The average number of syllables per sentence in the dataset is 50.89 for Japanese, whereas 36.18 per English and 40.40 per Korean. With these statistics, it can be deduced that Japanese necessitates approximately 41% more syllables than English and 26% more syllables than Korean to express an equivalent message. This limitation forces translators to often modify the meaning of original lyrics by adding, omitting, or even tweaking the message. However, translated lyrics still aim to capture the theme, mood, and spirit of the original lyrics [17]. Therefore, while original and translated lyrics need not be semantically identical, they still need to be semantically relevant [4, 18].

It is also crucial to preserve the frequency of phoneme repetition (e.g., rhyme) in translated lyrics, particularly when the music demands it [17]. For instance, some sections, such as choruses, require a substantial degree of phoneme repetition, while others do not. Moreover, due to the inherent connection between lyrics and music, lyrics must be arranged in a way that complements the music [19]. As a result, musically similar sections should maintain resembling linguistic features, including the choice of phonemes and the frequency of phoneme repetition [20].

3. DATASET

Although some websites provide user-translated multilingual lyrics, we found that most of them lack singability, as these translations were focused on delivering the meaning of the original lyrics rather than making them performable. While there are a few singable translations available, they are often not aligned on a line-by-line nor section-by-section basis due to the subjective nature of the lyric structure that there is no universal agreement on what to call a line and what to call a section. The absence of alignment makes it difficult to compare the original lyrics with their translated versions. To address this issue, we collected a set of singable lyrics, sourced from either official lyrics of commercial songs or user-translated ones found on YouTube, meticulously aligned on a line-by-line basis in English, Japanese, and Korean. This approach ensures that lyrics on the same line share the same melodies. Moreover, the dataset divides the lyrics into sections, allowing for section-by-section analysis. Alongside the lyrics, it

Section #	Line #	English (EN)	Japanese (JP)	Korean (KR)
1	1	Twinkle, twinkle, little star	きらきらひかる	반짝 반짝 작은별
	2	How I wonder what you are	おそらのほしよ	아름답게 비치네
	3	Up above the world so high	まばたきしては	서쪽 하늘에서도
	4	Like a diamond in the sky	みんなをみてる	동쪽 하늘에서도
	5	Twinkle, twinkle, little star	きらきらひかる	반짝 반짝 작은별
	6	How I wonder what you are	おそらのほしよ	아름답게 비치네
2	7	Twinkle, twinkle, little star	きらきらひかる	반짝 반짝 작은별
	8	How I wonder what you are	おそらのほしよ	아름답게 비치네
	9	When the blazing sun is gone	みんなのうたが	서쪽 하늘에서도
	10	When he nothing shines upon	とどくといいな	동쪽 하늘에서도
	11	Then you show your little light	きらきらひかる	반짝 반짝 작은별
	12	Twinkle, twinkle, all the night	おそらのほしよ	아름답게 비치네

Table 1. Sample data illustrating the original English lyrics of “Twinkle, Twinkle, Little Star” and their corresponding singable translations in Japanese and Korean, aligned on a line-by-line and section-by-section basis.

provides essential metadata such as genre, artist, original language, and the official status of lyrics. The dataset consists of 162 songs, each having lyrics in the three languages. It covers a diverse range of genres, including 109 K-pop, 23 animation music (e.g., Disney), 13 J-pop, 10 theatre music, and more. Table 1 shows sample data.

4. EVALUATING SINGABILITY

Our primary goal is to develop an evaluation framework that automatically assesses the quality of translated lyrics. One of the most important factors determining the quality is *singability*, defined as not only the ability of being sung, but also the suitability (easiness) of being sung [18]. To ensure such singability, we aim to provide metrics from three distinct perspectives by making sure that they i) maintain the song’s melodic integrity, ii) preserve the degree of phoneme repetition, and iii) consider the underlying musical structure.

To substantiate the reliability of our evaluation metrics, we conducted a comparative analysis of singable lyrics versus non-singable lyrics based on each proposed evaluation metric. In all our comparative analyses, we utilized our dataset for singable lyrics, where official lyrics served as both source and target lyrics, and unofficial functioned as only target lyrics. For non-singable lyrics, we obtained pairs of original singable (source) and human-translated non-singable (target) lyrics, aligned line-wise and section-wise, for 3,642 songs from <https://lyricstranslate.com/>.

Source	Target	Singable	Non-singable
English	Japanese	0.17 (80 songs)	0.74 (1401 songs)
	Korean	0.11 (80 songs)	0.48 (620 songs)
Japanese	English	0.16 (162 songs)	0.39 (589 songs)
	Korean	0.11 (162 songs)	0.31 (73 songs)
Korean	English	0.09 (161 songs)	0.20 (702 songs)
	Japanese	0.12 (161 songs)	0.52 (257 songs)

Table 2. The average line syllable count distance ($D_{iS_{syl}}$) between source and target languages for singable and non-singable lyrics.

Section	English	Japanese (English translation)	Korean (English translation)	rho
$E(A_1)$	Do you wanna build a snowman?	雪だるま作ろう (Let's build a snowman)	같이 눈사람 만들래? (Do you wanna build a snowman?)	
$J(A_1)$	Come on, let's go and play!	ドアを開けて (Please open the door)	제발 좀 나와봐 (Please come out)	0.85,
$K(A_1)$	I never see you anymore	一緒に遊ぼう (Let's play together)	언니를 만날 수 없어 (I can't meet you)	0.73,
$E(A_1)$	Come out the door	どうして (Why)	같이 놀자 (Let's play together)	0.77
$K(A_1)$	It's like you've gone away	出てこないの? (don't you come out?)	나 혼자 심심해 (I'm lonely alone)	
$E(B_1)$	We used to be best buddies	前は仲良く (We were close before)	그렇게 친했는데 (We were close before)	0.92,
$J(B_1)$	And now we're not	してたのに (We used to be)	이젠 아냐 (and we're not)	0.80,
$K(B_1)$	I wish you would tell me why!	なぜ会えないの (Why can't we meet each other?)	그 이유를 알고파 (I want to know the reason why)	0.91
$E(A_2)$	Do you wanna build a snowman?	雪だるま作ろう (Let's build a snowman)	같이 눈사람 만들래? (Do you wanna build a snowman?)	
$J(A_2)$	Or ride our bike around the halls?	自転車に乗ろう (Let's ride a bike)	아니면 자전거 탈래? (or do you wanna ride a bike?)	0.79,
$K(A_2)$	I think some company is overdue	ずっと一人でいると (When I'm alone all the time)	이제는 나도 지쳐 가나봐 (Seems I'm getting tired)	0.73,
$E(A_2)$	I've started talking to the pictures on the walls!	壁の絵とおしゃべりしちゃう (I'm almost talking to the pictures on the walls)	벽에다 말을 하며 놓고 있잖아 (because I've started talking to the walls)	0.82
$J(A_2)$	It gets a little lonely	さびしい部屋で (In a lonely room)	사실은 조금 외로워 (In fact, I'm a little lonely)	0.90,
$K(A_2)$	All these empty rooms	柱時計 (the wall clock)	텅빈 방에선 (In empty rooms)	0.88,
$E(B_2)$	Just watching the hours tick by	見てたりするの (I look at or something)	시계소리만 들려 (All I can hear is the clock's ticking)	0.96

Table 3. Lyric excerpt from “Do You Want to Build a Snowman” from the animation “Frozen,” singable in all languages. Sections A_1 and A_2 form a musically similar pair, while B_1 and B_2 are also musically similar to each other. Each section is denoted as $E(A_1), \dots, E(B_2)$ in English, $J(A_1), \dots, J(B_2)$ in Japanese, and $K(A_1), \dots, K(B_2)$ in Korean.

4.1 Line Syllable Count Distance

It is crucial to preserve the syllable counts between the original and translated lyrics for each line as similar as possible in order to maintain the integrity of a song’s melody [21]. Therefore, it is unsurprising that our evaluation framework incorporates a metric to assess the differences in syllable counts. Let the line syllable counts for a pair of lyrics that consist of n lines, $\mathbf{X} = \{x_1, \dots, x_n\}$ and $\tilde{\mathbf{X}} = \{\tilde{x}_1, \dots, \tilde{x}_n\}$ be denoted as $\{syl(x_1), \dots, syl(x_n)\}$ and $\{syl(\tilde{x}_1), \dots, syl(\tilde{x}_n)\}$ where each element refers to the syllable count of each line. For instance, if the first line of the English lyrics \mathbf{X} is “Silent night holy night” and the corresponding line in the Korean lyrics $\tilde{\mathbf{X}}$ is “Goyohanbam-georukhanbam (고요한밤 거룩한밤)”, the value of $syl(x_1)$ is 6 and $syl(\tilde{x}_1)$ is 8. We define the **line syllable count distance** between a pair of lyrics \mathbf{X} and $\tilde{\mathbf{X}}$ ($Dis_{syl}(\mathbf{X}, \tilde{\mathbf{X}})$) in order to evaluate the disparities in syllable counts, as follows.

$$Dis_{syl}(\mathbf{X}, \tilde{\mathbf{X}}) = \frac{1}{2n} \sum_{i=1}^n \left(\frac{|syl(x_i) - syl(\tilde{x}_i)|}{syl(x_i)} + \frac{|syl(x_i) - syl(\tilde{x}_i)|}{syl(\tilde{x}_i)} \right) \quad (1)$$

We compare the line syllable count distance of singable and non-singable lyrics. As shown in Table 2, non-singable lyrics display a considerably greater $Dis_{syl}(\mathbf{X}, \tilde{\mathbf{X}})$ compared to singable lyrics due to the varying syllable count requirements across languages.

4.2 Phoneme Repetition Similarity

Rhyme, defined as the repetition of a vowel sound and any subsequent sounds [22], has historically been a fundamental element in the realm of poetry, including in Western languages like English. However, the concept of rhyme has not been as prevalent in Japanese or Korean poetry [23]. In fact, traditional Korean poetry did not incorporate this concept [24]. Despite the increasing tendency to adopt the concept of rhyme in Japanese and Korean lyrics due to intercultural exchanges, we observed that lyrics in these languages often rely more on repeating grammatical elements. For example, in section A_1 of Table 3, the Japanese pair “tsukurou (作ろう, Let’s build)” and “asobou (遊ぼう, Let’s play)” generates a sense of rhyme because both end

with the same conjugation “ou” meaning “let’s”. Similarly, in Section A_2 , the Korean pair “mandeullae (만들래, Do you wanna build)” and “tallae (탈래, Do you wanna ride)” creates a sense of repetition because both end with “llae” meaning “Do you wanna”. Another example is the repetition of particles at the end of sentences, such as “yo (よ)” and “no (の)” in Japanese and “yo (요)” and “da (다)” in Korean, which convey cultural nuances related to formality. We therefore propose that English, Japanese, and Korean share common ground in adopting phoneme repetition for poetic expression. However, as such repetition is not necessarily called rhyme in Japanese and Korean, we will refrain from using the term “rhyme” and instead employ the term “phoneme repetition.”

We noticed that each section’s degree of phoneme repetition remains consistent across different languages when the lyrics are singable. For example, in Table 3, the first section of the original English lyrics ($E(A_1)$) displays a strong degree of phoneme repetition, with three rhyming pairs: “come-come”, “play-away”, and “anymore-door” (In this paper, we denote a section as an uppercase with a number and a line as a lower case with a number). Similarly, both the Japanese and Korean translations ($J(A_1)$, $K(A_1)$) also exhibit a substantial degree of phoneme repetition, featuring three pairs of repeated phonemes in each: “doa (ドア)”-“dou (どう)”, “tsukurou (作ろう)”-“asobou (遊ぼう)”, “akete (開けて)”-“shite (して)” in Japanese, and “gachi (같이)”-“gachi (같이)”, “mandeul (만들)”-“eonnireul (언니를)”, “mandeullae (만들래)”-“simsimhae (심심해)” in Korean. However, we realized that it is not fair to directly compare the number of phoneme repetitions when attempting to quantify the degree of phoneme repetition as each language has a different number of vowels and consonants: English has 15 vowels and 24 consonants, whereas Japanese has 5 and 15 and Korean has 21 and 19. Hence, in an attempt to minimize the differences in the number of phonemes, we treated acoustically similar vowels as the same vowel in English, such as ‘IH’-‘IY’, ‘UH’-‘UW’, or ‘EH’-‘AE’ (e.g., ‘mass’ and ‘mess’) [25] because they can still form slant rhymes [26]. Conversely, we considered ‘A’-‘YA’, ‘O’-‘YO’, and ‘U’-‘YU’ as sep-

Source	Target	Singable	Non-singable (Human)	Non-singable (Machine)
English	Japanese	0.14	0.26	0.30
	Korean	0.10	0.15	0.15
Japanese	English	0.13	0.20	0.18
	Korean	0.11	0.13	0.14
Korean	English	0.10	0.10	0.10
	Japanese	0.11	0.10	0.17

Table 5. The average musical structure distance (Dis_{mus}) of singable lyrics and non-singable lyrics.

between lyrics in different languages \mathbf{X} and $\tilde{\mathbf{X}}$, each consisting of m sections, $Dis_{mus}(\mathbf{X}, \tilde{\mathbf{X}})$, is defined as follows:

$$Dis_{mus}(\mathbf{X}, \tilde{\mathbf{X}}) = \frac{1}{m^2} \sqrt{\sum_{i,j=1}^m (diss(X_i, X_j) - diss(\tilde{X}_i, \tilde{X}_j))^2} \quad (5)$$

In Table 5, we provide a summary of the average musical structure distance for singable lyrics, human-translated non-singable lyrics, and machine-translated non-singable lyrics generated by automatically translating official singable lyrics from 80 English, 162 Japanese, and 161 Korean songs using Google Translator. Our findings show that singable lyrics exhibit the lowest Dis_{mus} values, while machine-translated non-singable lyrics display the highest, suggesting that machine-translated ones lack structural coherence the most. As human-translated non-singable lyrics maintain structural coherence in aspects such as word choice and nuances, they demonstrate lower distances than machine-translated counterparts.

5. EVALUATING SEMANTICS

Semantic relatedness to the original lyrics is by no means less fundamental than syllable counts, phoneme repetition, and structural factors [18]. We therefore introduce a fourth metric, semantic similarity, to ensure the semantic relevance of translated lyrics to the original.

5.1 Semantic Similarity

To numerically assess the semantic textual similarity (sts) between a pair of lyrics, we first obtained the contextual embeddings of each text from lyrics using a pre-trained Sentence BERT model¹ [31] and then calculated the cosine similarity between the embeddings. As this model was trained for English, the Japanese and Korean lyrics were automatically translated using Google Translator before obtaining the embeddings.

We started by examining hierarchical semantic similarity using cross-scape plots [32], as shown in Figure 2. Given a pair of lyrics $\mathbf{X} = \{x_1, \dots, x_n\}$ and $\tilde{\mathbf{X}} = \{\tilde{x}_1, \dots, \tilde{x}_n\}$ with n lines each, the first (leftmost) block of the lowest line represents the semantic textual similarity between x_1 and \tilde{x}_1 (denoted as $sts(x_1, \tilde{x}_1)$) while the last (rightmost) block signifies $sts(x_n, \tilde{x}_n)$. The first (leftmost) block of the second-lowest line denotes $sts(x_1+x_2, \tilde{x}_1+\tilde{x}_2)$, and the second block corresponds to $sts(x_2+x_3, \tilde{x}_2+\tilde{x}_3)$. Lastly,

¹ We used all-MiniLM-L6-v2 [30].

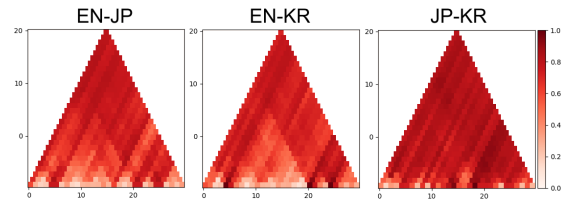


Figure 2. Semantic similarity cross-scape plot for the J-pop song “A Thousand Winds” between English and Japanese (Left), English and Korean (Middle), and Japanese and Korean (Right). Any value less than 0 was regarded as 0.

Line #	English	Japanese (English translation)	sts
1	please do not stand at my grave and weep.	私のお墓の前で (in front of my grave)	0.56
2	I am not there. I do not sleep	泣かないでください (please stop crying)	0.27
1, 2	please do not stand at my grave and weep. I am not there, I do not sleep	私のお墓の前で 泣かないでください (in front of my grave. please stop crying.)	0.76

Table 6. Semantic textual similarity (sts) between English and Japanese versions of “A Thousand Winds”.

the highest block (line) represents the similarity between the entire lyrics, $sts(x_1 + \dots + x_n, \tilde{x}_1 + \dots + \tilde{x}_n)$.

In each plot of Figure 2, there are semantic disparities at lower levels, but similarities increase at higher (broader) levels. We have two explanations for this. First, the number of musical notes within a single lyric line may be adequate to deliver a specific message in one language but insufficient in another language. Therefore, it is common for a message conveyed in one line in one language to span two lines in another language. As an example, we provide Table 6, which presents the semantic textual similarity (sts) between Japanese and English lyrics of the J-pop song “A Thousand Winds (千の風になつて)”. As demonstrated in the table, the similarity between Japanese and English at a broader level ($sts(x_1+x_2, \tilde{x}_1+\tilde{x}_2)$) can be higher than at the line level ($sts(x_1, \tilde{x}_1)$, $sts(x_2, \tilde{x}_2)$) because Japanese generally requires more syllables than English and it often takes two lines in Japanese to express a single-line message in English. Second, the semantic similarities at broader levels can be higher because of grammatical/linguistic differences. Each language has its own natural word order patterns. For example, in the phrase “I’m going to travel to find the gold,” it is natural in English to mention “I’m going to travel” before “to find the gold.” However, in Japanese and Korean, expressing “to find the gold (金を探しに, 금을 찾으러)” before “I’m going to travel (旅に出る, 떠난다)” is a more typical and natural construction. Table 7 shows that these differences between languages make line-level semantic assessments insufficient. Since lines 1, 2, and 3 in the English version correspond to lines 3, 1, and 2 respectively in the Japanese version, these pairs exhibit low semantic similarities at the line level ($sts(x_1, \tilde{x}_1)$, $sts(x_2, \tilde{x}_2)$, $sts(x_3, \tilde{x}_3)$), while demonstrating higher similarity when considered as a whole ($sts(x_1+x_2+x_3, \tilde{x}_1+\tilde{x}_2+\tilde{x}_3)$).

Considering these factors, it becomes evident that

Line #	English	Japanese (English translation)	<i>sts</i>
1	Dare to try and reach out for heaven	望むように生きるなら (If you want to become what you're meant to be)	0.22
2	You must become what you're meant to be	星からの金を求め (to find the gold from stars)	0.27
3	And bring the gold of heaven to the world	一人旅に出るのよ (Dare to embark on a solo journey)	0.13
1-3	Dare to try and reach out for heaven You must become what you're meant to be And bring the gold of heaven to the world	望むように生きるなら星からの金を求め一人旅に出るのよ (Dare to embark on a solo journey if you want to become what you're meant to be to find the gold from stars)	0.53

Table 7. Semantic textual similarity (*sts*) between English and Japanese versions of “Gold von den Sternen”.

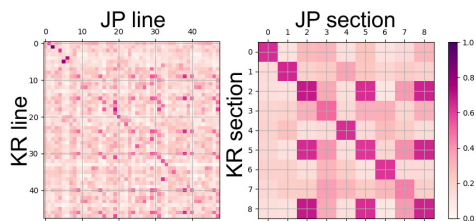


Figure 3. The line-wise (Left) and section-wise (Right) semantic similarity matrices between Japanese and Korean versions of “Wie wird man seinen Schatten los?”

singable lyric translations do not prioritize line-wise semantic similarity. Rather, we observed that singable translations aim to preserve semantic connections at the section level since the organization of the lyric storyline follows a section-wise approach. To illustrate this, we present Figure 3, which displays both line-wise and section-wise semantic similarity matrices for the Japanese and Korean versions of “How do you get rid of your shadow? (Wie wird man seinen Schatten los?)” from the German musical “Mozart!”. As shown in the Figure, the section-wise matrix represents the semantic relatedness more clearly than the line-wise matrix.

Therefore, we propose assessing section-wise semantic relatedness for evaluating singable lyric translation. To achieve this, we define the **semantic similarity** between a pair of lyrics $\mathbf{X} = \{X_1, \dots, X_m\}$ and $\tilde{\mathbf{X}} = \{\tilde{X}_1, \dots, \tilde{X}_m\}$, consisting of m sections and $n = n(X_1) + \dots + n(X_m)$ lines, where $n(X_i)$ denotes the number of lines in the i -th section, as follows:

$$Sim_{sem}(\mathbf{X}, \tilde{\mathbf{X}}) = \sum_{i=1}^m \left(\frac{n(X_i)}{n} sts(X_i, \tilde{X}_i) \right). \quad (6)$$

Table 8 compares singable and non-singable lyrics in terms of line-wise semantic similarities ($\frac{1}{n} \sum_{i=1}^n sts(x_i, \tilde{x}_i)$) and section-wise similarities, using our proposed metric (Sim_{sem}). The table reveals that non-singable translations exhibit high semantic similarity for both line-wise and section-wise measures, with similar values for each. In contrast, singable translations show low line-wise similarity, as expected, since they do not prioritize line-wise semantic similarity. However, when evaluated using section-wise similarity, they display a level of similarity comparable to that between “Machine learning is so easy” and “Deep learning is so straightforward”, which is 0.623 when measured with the same pre-trained model [30].

Source	Target	Singable		Non-singable	
		line	section	line	section
English	Japanese	0.40	0.54	0.64	0.74
	Korean	0.42	0.55	0.70	0.76
Japanese	English	0.47	0.59	0.66	0.72
	Korean	0.52	0.61	0.77	0.79
Korean	English	0.53	0.63	0.78	0.81
	Japanese	0.52	0.61	0.73	0.75

Table 8. The average line-wise semantic similarity and section-wise semantic similarity (Sim_{sem}) of singable and non-singable lyrics.

6. DISCUSSIONS AND CONCLUSIONS

In this paper, we introduced a computational evaluation framework for singable lyric translation, grounded in the musical, linguistic, and cultural understanding of lyrics, comprised of four evaluation metrics, line syllable count distance (Dis_{syl}), phoneme repetition similarity (Sim_{pho}), musical structure distance (Dis_{mus}), and semantic similarity (Sim_{sem}). These metrics are designed to ensure that the translated lyrics maintain the integrity of melodies, degree of phoneme repetition, structural coherence, and semantics of the original lyrics. Our framework is automated, guaranteeing objectivity and efficiency in terms of time and cost. We showed the efficacy of our evaluation metrics by offering comparative statistics between singable and non-singable lyrics. In addition, our analysis revealed that the degree of phoneme repetition in the original lyrics is frequently mirrored in the translated lyrics, musically similar sections tend to share the same phonemes and display comparable degrees of phoneme repetition, and section-wise analysis is better suited for evaluating semantic similarity for lyric translation than line-wise analysis.

Nonetheless, there remains room for improvement. Although we have assembled a singable lyrics dataset, aligned across English, Japanese, and Korean, our dataset has some limitations; it lacks musical information and its volume is limited. As a result, we have not been able to incorporate musical notes into our experiment or conduct comparative studies across various genres. We recognize that an ideal lyric translation evaluation system should take into account the relationship between musical notes and phonemes, as well as adapt to different genres. Moreover, although we have endeavored to incorporate cultural understandings of poetry in different languages, we acknowledge the need for deeper cultural considerations. For example, we noticed that cultural similarities might have an impact on the extent of semantic similarities. This is demonstrated in an English translation of “MIC Drop”, a K-pop song by BTS originally written in Korean, made by YouTuber Iris Phuong. The translated singable lyrics do not include a translation of the term “hyodo (효도, taking care of parents)” as there is no direct equivalent in English, while the Japanese version of the song effortlessly conveys the concept as “kougou (孝行)”. In the future, we aim to expand our dataset to contribute more to lyric translation studies and to further explore the impact of genre and cultural influences on lyric translation.

7. ACKNOWLEDGMENTS

This work was supported in part by JST CREST Grant Number JPMJCR20D4.

8. REFERENCES

- [1] M. Mateo, "Music and translation," *Handbook of translation studies*, vol. 3, pp. 115–121, 2012.
- [2] Ş. Susam-Sarajeva, "Translation and music: Changing perspectives, frameworks and significance," *The Translator*, vol. 14, no. 2, pp. 187–200, 2008.
- [3] S. Spaeth, "Translating to music," *The Musical Quarterly*, vol. 1, no. 2, pp. 291–298, 1915.
- [4] P. Low, "Translating songs that rhyme," *Perspectives: Studies in translatology*, vol. 16, no. 1-2, pp. 1–20, 2008.
- [5] —, "The pentathlon approach to translating songs," in *Song and significance*. Brill, 2005, pp. 185–212.
- [6] F. Guo, C. Zhang, Z. Zhang, Q. He, K. Zhang, J. Xie, and J. Boyd-Graber, "Automatic song translation for tonal languages," in *Findings of the Association for Computational Linguistics (ACL)*. Dublin, Ireland: Association for Computational Linguistics, May 2022, pp. 729–743. [Online]. Available: <https://aclanthology.org/2022.findings-acl.60>
- [7] L. Ou, X. Ma, M.-Y. Kan, and Y. Wang, "Songs across borders: Singable and controllable neural lyric translation," *arXiv preprint arXiv:2305.16816*, 2023.
- [8] C. Li, K. Fan, J. Bu, B. Chen, Z. Huang, and Z. Yu, "Translate the beauty in songs: Jointly learning to align melody and translate lyrics," *arXiv preprint arXiv:2303.15705*, 2023.
- [9] K. Watanabe and M. Goto, "Lyrics information processing: Analysis, generation, and applications," in *Proceedings of the 1st Workshop on NLP for Music and Audio (NLP4MusA)*, 2020, pp. 6–12.
- [10] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, "BLEU: A method for automatic evaluation of machine translation," in *Proceedings of the 40th annual meeting of the Association for Computational Linguistics (ACL)*, 2002, pp. 311–318.
- [11] C.-Y. Lin, "ROUGE: A package for automatic evaluation of summaries," in *Text summarization branches out*, 2004, pp. 74–81.
- [12] S. Banerjee and A. Lavie, "METEOR: An automatic metric for MT evaluation with improved correlation with human judgments," in *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, 2005.
- [13] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, and Y. Artzi, "BERTScore: Evaluating text generation with bert," in *International Conference on Learning Representations (ICLR)*, 2019.
- [14] I. Marc, "Travelling songs: On popular music transfer and translation," *IASPM Journal*, vol. 5, no. 2, pp. 3–21, 2015.
- [15] E. C. Hui-tung, "Translation of songs," *An Encyclopedia of Practical Translation and Interpreting*, p. 351, 2019.
- [16] Y. Yang, Y. Zhang, C. Tar, and J. Baldridge, "PAWS-X: A cross-lingual adversarial dataset for paraphrase identification," 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)*, 2019, pp. 3687–3692.
- [17] H. S. Drinker, "On translating vocal texts," *The Musical Quarterly*, vol. 36, no. 2, pp. 225–240, 1950.
- [18] J. Franzon, "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, pp. 333–346, 2015.
- [19] J. P. G. Mahedero, A. Martínez, P. Cano, M. Koppenberger, and F. Gouyon, "Natural language processing of lyrics," in *Proceedings of the 13th Annual ACM International Conference on Multimedia*, ser. MULTIMEDIA '05. Association for Computing Machinery, 2005, pp. 475–478.
- [20] K. Watanabe and M. Goto, "A chorus-section detection method for lyrics text." in *Proceedings of the 21th International Conference on Music Information Retrieval (ISMIR)*, 2020, pp. 351–359.
- [21] R. Apter and M. Herman, "Translating art songs for performance: Rachmaninoff's six choral songs," *Translation Review*, vol. 84, no. 1, pp. 27–42, 2012.
- [22] A. Bain, "English composition and rhetoric,(2nd american ed.)," *New York: D. Appleton and Company*, 1867.
- [23] N. Manabe, "Globalization and Japanese creativity: Adaptations of Japanese language to rap," *Ethnomusicology*, vol. 50, no. 1, pp. 1–36, 2006.
- [24] Y. B. Yoon and B. L. Derwing, "A language without a rhyme: Syllable structure experiments in Korean," *Canadian Journal of Linguistics/Revue canadienne de linguistique*, vol. 46, no. 3-4, pp. 187–237, 2001.
- [25] A. K. Syrdal and H. S. Gopal, "A perceptual model of vowel recognition based on the auditory representation of American English vowels," *The Journal of the Acoustical Society of America*, vol. 79, no. 4, pp. 1086–1100, 1986.

- [26] K. Hanson, “Formal variation in the rhymes of Robert Pinsky’s the inferno of dante,” *Language and Literature*, vol. 12, no. 4, pp. 309–337, 2003.
- [27] C. Ito, Y. Kang, and M. Kenstowicz, “The adaptation of Japanese loanwords into Korean,” *MIT Working Papers in Linguistics*, vol. 52, no. 2006, pp. 65–104, 2006.
- [28] S.-E. Chang, “Enhancement effects of clear speech and word-initial position in Korean glides,” *The Journal of the Acoustical Society of America*, 2017.
- [29] J. Li, M. Galley, C. Brockett, J. Gao, and W. B. Dolan, “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 (NACCL) : Human Language Technologies*, 2016, pp. 110–119.
- [30] W. Wang, F. Wei, L. Dong, H. Bao, N. Yang, and M. Zhou, “MiniLM: Deep self-attention distillation for task-agnostic compression of pre-trained transformers,” *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, pp. 5776–5788, 2020.
- [31] N. Reimers and I. Gurevych, “Sentence-BERT: Sentence embeddings using Siamese BERT-networks,” 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)*, 2019, pp. 3982–3992.
- [32] S. Park, T. Kwon, J. Lee, J. Kim, and J. Nam, “A cross-scape plot representation for visualizing symbolic melodic similarity,” in *Proceedings of the 20th International Conference on Music Information Retrieval (ISMIR)*, 2019, pp. 423–430.