

# LYRICWHIZ: ROBUST MULTILINGUAL ZERO-SHOT LYRICS TRANSCRIPTION BY WHISPERING TO CHATGPT


Le Zhuo<sup>1</sup> Ruibin Yuan<sup>2,3</sup> Jiahao Pan<sup>4</sup> Yinghao Ma<sup>5</sup> Yizhi Li<sup>6</sup> Ge Zhang<sup>2,7</sup> Si Liu<sup>1</sup>  
Roger Dannenberg<sup>3</sup> Jie Fu<sup>2</sup> Chenghua Lin<sup>6</sup> Emmanouil Benetos<sup>5</sup> Wenhui Chen<sup>7</sup> Wei Xue<sup>4</sup> Yike Guo<sup>4</sup>  
<sup>1</sup> Beihang University <sup>2</sup> Beijing Academy of Artificial Intelligence <sup>3</sup> Carnegie Mellon University  
<sup>4</sup> Hong Kong University of Science and Technology <sup>5</sup> Queen Mary University of London  
<sup>6</sup> University of Sheffield <sup>7</sup> University of Waterloo  
zhuole1025@gmail.com, ruibiny@andrew.cmu.edu, fujie@baai.ac.cn

## ABSTRACT

We introduce LyricWhiz, a robust, multilingual, and zero-shot automatic lyrics transcription method achieving state-of-the-art performance on various lyrics transcription datasets, even in challenging genres such as rock and metal. Our novel, training-free approach utilizes Whisper, a weakly supervised robust speech recognition model, and GPT-4, today’s most performant chat-based large language model. In the proposed method, Whisper functions as the “ear” by transcribing the audio, while GPT-4 serves as the “brain,” acting as an annotator with a strong performance for contextualized output selection and correction. Our experiments show that LyricWhiz significantly reduces Word Error Rate compared to existing methods in English and can effectively transcribe lyrics across multiple languages. Furthermore, we use LyricWhiz to create the first publicly available, large-scale, multilingual lyrics transcription dataset with a CC-BY-NC-SA copyright license, based on MTG-Jamendo, and offer a human-annotated subset for noise level estimation and evaluation. We anticipate that our proposed method and dataset will advance the development of multilingual lyrics transcription, a challenging and emerging task.

## 1. INTRODUCTION

Automatic lyrics transcription (ALT) is a crucial task in music information retrieval (MIR) that involves converting an audio recording into a textual representation of the lyrics sung in the recording. The importance of this task stems from the fact that lyrics are a fundamental aspect of many music genres and are often the main way in which listeners engage with and interpret a song’s meaning. Additionally, ALT has numerous applications in the music industry, such as enabling better cataloging [1], music searching [2, 3], music recommendation [4], as well as



**Figure 1.** Concept illustration of the working LyricWhiz, where user prompts the two advanced models, Whisper and ChatGPT, to perform automatic lyrics transcription.

facilitating the creation of karaoke tracks and lyric videos. Moreover, ALT can assist in various music-related research tasks, including sentiment analysis [5], music genre classification [1], lyrics generation, which is further used for music generation [6], security review, and music copyright protection. Thus, accurate and efficient ALT is essential for advanced MIR and the development of new music-related applications.

However, to date, no sufficiently robust and accurate ALT system has been developed. Even major commercial music streaming platforms still rely heavily on manually-annotated lyrics, incurring high costs. One key reason is the challenging nature of lyrics transcription. The diversity of singing styles and skills leads to varied timbres of the same pronunciation. Moreover, the phonemes in singing may be pronounced in vastly different ways, such as longer duration, tone changes, or even vowel substitutions, to accommodate the melody. Lastly, the inclusion of various music accompaniments across different genres makes it challenging to distinguish the vocal signals from other sounds. To surmount these challenges, a more robust ALT system is necessary, capable of outperforming existing models in diverse scenarios, including the transcription of multilingual lyrics.

Another significant factor hindering the progress of



© L Zhuo, R Yuan, and J Pan. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). **Attribution:** L Zhuo, R Yuan, and J Pan, “LyricWhiz: Robust Multilingual Zero-shot Lyrics Transcription by Whispering to ChatGPT”, in *Proc. of the 24th Int. Society for Music Information Retrieval Conf.*, Milan, Italy, 2023.

ALT systems is the absence of large-scale singing datasets. Currently, only two relatively sizable datasets [7, 8] exist for ALT systems. However, all existing datasets are in English, with no multilingual datasets available. Besides, these datasets often have stringent copyright licensing restrictions, which significantly hampers their utilization by researchers. Consequently, developing a more comprehensive and representative dataset, encompassing multiple languages and without copyright issues, is essential for supporting the creation of a robust and accurate system.

In this paper, we present LyricWhiz, a novel method for automatic lyrics transcription. LyricWhiz surpasses existing methods on various ALT datasets, resulting in a significant reduction in WER for English lyrics and providing accurate transcription results across multiple languages. Our system is robust, multilingual, and training-free. To achieve these results, we combined two powerful models from their respective domains as shown in Figure 1: Whisper, a weakly supervised speech transcription model, and GPT-4, a large language model (LLM) from the ChatGPT family. Whisper acts as the “ear” while GPT-4 serves as the “brain” by providing contextualized output selection and correction with strong performance [9]. We further use LyricWhiz to build a multilingual lyrics dataset, named MulJam, which is the first large-scale, multilingual lyrics transcription dataset without copyright-related issues.

The contributions of our work are as follows:

- We propose a novel, robust, training-free ALT method, LyricWhiz, which significantly reduces WER on various ALT benchmark datasets, including Jamendo, Hansen, and MUSDB18, and is close to the in-domain state-of-the-art system on DSing.
- We introduce the first ALT system that can perform zero-shot, multilingual, long-form ALT by integrating a large speech transcription model and an LLM for contextualized post-processing.
- We create the first publicly-available, large-scale, multilingual lyrics transcription dataset with a clear copyright statement which eliminates further reviewing of the users and facilitates public usage. We provide a human-annotated subset to estimate noise levels and evaluate multilingual ALT performance.

## 2. RELATED WORK

### 2.1 Automatic Lyrics Transcription

Automatic lyrics transcription (ALT) is an essential task in music information retrieval and analysis, aiming to recognize lyrics from singing voices. It remains challenging due to facts such as the sparsity of training data and the unique acoustic characteristics of the singing voice that differ from normal speech. Traditional methods treat ALT in the automatic speech recognition (ASR) framework, which generally utilizes a hybrid of language model and acoustic model, e.g., HMM-GMM. Music-related characteristics have been used to further address these challenges [11–13].

Despite integrating domain-specific music priors into system designs, the data scarcity issue persists. Recently, some researchers have constructed datasets for end-to-end learning, which greatly advances ALT, but most datasets are either noisy (DALI [7, 14], Hansen [15], DAMP-MVP<sup>1</sup>); not large (Vocadito [16]); or not diverse in terms of genre and language (MUSDB18 [17], DSing [8]).

Recent rapid progress in ASR has greatly benefited ALT. Some work focuses on applying the ASR model architectures [18–20], such as the Transformers, to ALT, and other work leverages the vast amount of public annotated ASR datasets [19–21] to bridge between the speech and music data. For the first time, a recent study [22] transferred a large-scale self-supervised pre-trained ASR model, mus2vec 2.0, to the singing domain, and exhibited superior performance on multiple benchmark datasets. Nevertheless, this approach consists of pre-training, fine-tuning, and transfer learning phases, thereby remaining relatively complicated and still requiring singing datasets.

### 2.2 Weakly Supervised Automatic Speech Recognition

The paradigm of large-scale unsupervised pretraining and non-large annotated dataset finetuning has dominated end-to-end ASR research [23]. Well-known pretrained ASR models include contrastive learning based Wav2vec [24], Wav2vec 2.0 [25], HuBert [26], WavLM [27], Whisper [28], and Vall-E [29], which have performed impressively in various downstream tasks, including ASR and speech synthesis. Among them, Whisper has been most recognized for its ASR robustness across different datasets and its multilingual and multitasking capabilities, making Whisper potentially applicable to music tasks. Besides, specifically for ALT, pre-trained musical audio models including JukeBox [6], MusicLM [30], MULE [31], SingSong [32], music2vec [33], and MERT [34], may also contribute to achieving strong performance.

### 2.3 Chat-based Large Language Models

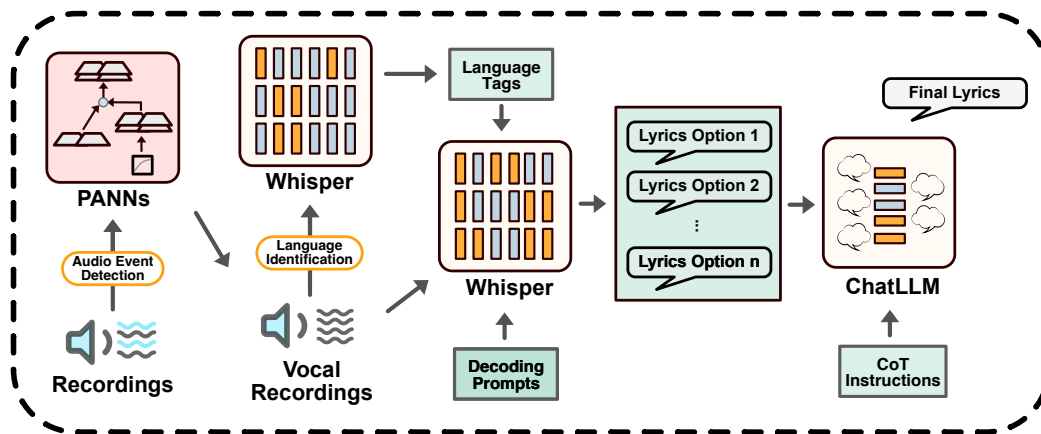
ChatGPT<sup>2</sup>, a chat-based large language model (LLM), has found broad application in optimizing workflows across a variety of domains, including multimodal intelligence [35, 36]. Recent breaking AutoGPT<sup>3</sup> is even recognized as an embryonic form of artificial general intelligence. Inspired by these developments, LyricWhiz collaborates with both Whisper [28] and ChatGPT to optimize the workflow of ALT. Prompt engineering is known to be important to navigate LLMs to perform better [37]. LyricWhiz mainly adopts three primary strategies:

a) As shown in [38, 39], a well-formalized task description prompt can effectively improve ChatGPT’s performance on downstream tasks with strict format requirements. We follow this empirical observation to strictly formalize the expected format of ALT post-processing outputs. We also refer to the prompt pattern catalog in [39] for an intuitive understanding of prompt engineering.

<sup>1</sup> <https://zenodo.org/record/2747436#.ZDqBQOzML0o>

<sup>2</sup> <https://openai.com/blog/chatgpt>

<sup>3</sup> <https://github.com/Torantulino/Auto-GPT>



**Figure 2.** Framework of the proposed LyricWhiz. In the first stage, we employ PANNs [10], to detect audio events and filter out non-vocal recordings. In the second stage, we utilize the language identification module in Whisper to predict input audio language. We then construct language-specific prompts for Whisper and transcribe input audio multiple times. In the final stage, we request ChatGPT with CoT instructions to ensemble multiple predictions and generate the final lyrics.

b) Inspired by [40, 41], LyricWhiz utilizes prompt augmentation to ask ChatGPT to analyze the prompt and input lyrics, in order to select the most accurate prediction from multiple Whisper trials, which is done in the first phase as illustrated in Section 3.2. [41] designs a gradient-guided strategy to select prompts. By contrast, we simply feed ChatGPT with an instruction to select prompts for itself.

c) The importance of a well-designed CoT [42], which effectively divides a complicated task into several phases and designs specific prompts for each phase, is widely acknowledged for enhancing LLM performance. We also propose a concise CoT strategy, depicted in Section 3.2.

### 3. METHODOLOGY

The overall framework of our method is presented in Figure 2. This section will provide an in-depth analysis of the design of the Whisper and ChatGPT components, and our multilingual dataset.

#### 3.1 Whisper as Zero-shot Lyrics Transcriber

In the Whisper [28] paper, the authors scaled the weakly supervised ASR to 680,000 hours of labeled audio data, which covers 96 languages and includes both multilingual and multitask training. This approach demonstrates high-quality results without the need for fine-tuning, self-supervision, or self-training techniques. By leveraging weak supervision and large-scale training, Whisper generalizes well to standard benchmarks and achieves robust speech recognition in various downstream tasks.

Motivated by this, we discovered that the weakly supervised Whisper model, trained on speech data, also excels in lyrics transcription within the music domain. We directly apply Whisper to transcribe lyrics of music from various genres, including pop, folk, rock, and rap, and find that the model consistently achieves accurate transcription results. The model excels at long-form transcription and is robust to different song styles, even for challenging genres such as rock and electronic music, where Whisper still provides

reasonable results. We further test Whisper on multiple benchmark datasets for lyric transcription. The results indicate that Whisper, without any training or fine-tuning, can achieve or surpass SOTA performance across multiple lyric transcription datasets.

Upon analyzing the transcription results from Whisper, we observed that the model occasionally outputs content unrelated to lyrics, such as music descriptions, emojis, website watermarks, and YouTube advertisements. We attribute this to the weakly supervised training of Whisper on large-scale noisy speech datasets. To address this issue, we utilize the input prompt designed in Whisper as a prefix prompt to guide it toward the lyric transcription task. Unlike prompt designing philosophy in other large language models, Whisper’s prefix prompt does not work well with explicit task instructions and has difficulty understanding lengthy explanations. In practice, we notice that using the simplest prompt, “lyrics:”, effectively prevents the model from outputting descriptions of the music in most cases, resulting in a significant improvement in transcription results. Therefore, in the following sections, this prompt is consistently used for Whisper’s transcription input.

Additionally, we apply post-processing tricks to Whisper’s output, utilizing the model’s predicted no-speech probability to handle situations where predictions are made despite the absence of vocals in the song. Specifically, we drop predicted lines of lyrics with a no speech probability greater than 0.9. This effectively filters out watermarks and advertisements, further enhancing the transcription results.

#### 3.2 ChatGPT as Effective Lyrics Post-processor

Although we addressed some issues with Whisper’s predictions through prompt design and post-processing, we still cannot avoid transcription translation errors, as well as grammatical and syntactical errors. Furthermore, due to the inherently stochastic nature of temperature scheduling in Whisper, the transcription predictions vary with each run, leading to fluctuations in evaluation metrics. To reduce this variance and enhance overall accuracy, we gen-

GPT-4 Instruction Prompt
<p>Task: As a GPT-4 based lyrics transcription post-processor, your task is to analyze multiple ASR model-generated versions of a song’s lyrics and determine the most accurate version closest to the true lyrics. <b>Also filter out invalid lyrics when all predictions are nonsense.</b></p> <p>Input: The input is in JSON format:  <pre>{“prediction_1”: “line1;line2;...”, ...}</pre></p> <p>Output: Your output must be strictly in readable JSON format without any extra text:  <pre>{   “reasons”: “reason1;reason2;...”,   “closest_prediction”: &lt;key_of_prediction&gt;   “output”: “line1;line2...” }</pre></p> <p>Requirements: For the "reasons" field, you have to provide a reason for the choice of the "closest_prediction" field. For the "closest_prediction" field, choose the prediction key that is closest to the true lyrics. <b>Only when all predictions greatly differ from each other or are completely nonsense or meaningless, which means that none of the predictions is valid, fill in "None" in this field.</b> For the "output" field, you need to output the final lyrics of closest_prediction. <b>If the "closest_prediction" field is "None", you should also output "None" in this field. The language of the input lyrics is English.</b></p>

**Table 1.** Instruction prompt for GPT-4 contextualized post-processing. We decompose this task into three consecutive phases, inspired by Chain-of-Thought prompting. Note that lines in blue indicate additional prompts used exclusively for multilingual dataset construction.

erate 3 to 5 predictions for each input music under identical settings and employ ChatGPT as an expert in lyrics to ensemble these multiple predictions.

The crux of the problem lies in designing an effective prompt for ChatGPT to accomplish the ensemble task reasonably. As shown in Table 1, we first assign ChatGPT the role of a transcription post-processor, indicating that its task is to analyze multiple lyric transcription results and select the one it deems most accurate. We then stipulate that both input and output should be in JSON format to facilitate structured processing and provide detailed descriptions for each output field.

Drawing on the Chain-of-Thought in large language models for reasoning, we devised a concise thought chain for ChatGPT that decomposes lyrics post-processing into three consecutive phases. This involves first having ChatGPT analyze multiple lyric inputs and provide reasons for selection, then making a choice, and finally outputting the chosen lyric prediction. We test this approach using GPT-3.5 and the newly released GPT-4. The results demonstrate that using the analysis-selection-prediction prompt for ChatGPT’s inference effectively enhances the final transcription results, with GPT-4 exhibiting a noticeably superior performance compared to GPT-3.5.

### 3.3 Multilingual Lyrics Transcription Dataset

Building upon the exceptional performance of the proposed framework in lyric transcription tasks, we further extend it to the challenging task of multilingual lyric transcription, introducing the first large-scale, weakly supervised, and copyright-free multilingual lyric transcription dataset. We utilize the publicly available MTG-Jamendo

Dataset	Languages	Songs	Lines	Duraion
DSing [8]	1 (en)	4,324	81,092	149.1h
MUSDB18 [17]	1 (en)	82	2,289	4.6h
DALI-train [14]	1 (en)	3,913	180,034	208.6h
DALI-full [14]	30*	5,358*	-	-
MulJam (Ours)	6	6,031	182,429	381.9h

**Table 2.** Comparison between different lyrics transcription datasets. Our model operates with a longer window (~30s), resulting in fewer lines compared to other datasets.

dataset for music classification, which comprises 55,000 full audio tracks, 195 tags, and music in various languages.

Since the MTG dataset contains a considerable proportion of non-vocal music, we first employ PANNs [10], a large-scale pre-trained audio pattern recognition model, to detect audio events and filter out non-vocal music with vocal-related tag probabilities below a predefined threshold. This filtering method eliminates approximately 60% of the music, thereby substantially reducing the time and resources required for dataset construction. We then utilize Whisper to transcribe lyrics from the music.

As the music in the MTG dataset encompasses multiple languages, we first utilize the Language Identification module within Whisper to predict the language of input music. Based on the predicted language, we translate the prefix prompt “lyrics:” into the corresponding language for input, e.g., “paroles” in French, and “liedtext” in German. After obtaining the transcription results, we discard lyrics that are too short or too long. When ensembling the prediction results with ChatGPT, we also incorporate the language of lyrics as an input condition in the prompt. Given the prevalence of nonsensical content in the transcription results, we additionally require ChatGPT to evaluate the validity of the transcribed lyrics in the prompt. If all input lyrics are deemed nonsensical, e.g., all special Unicode characters, or extremely divergent, the transcription result for that piece of music is considered invalid and discarded.

To prepare the dataset for training, it is essential to conduct line-level annotation. Timestamps can be obtained from the output of Whisper by aligning the lyrics both before and after ChatGPT processing. For the alignment of strings, the Levenshtein distance [43] is employed. To exclude aligned lines of lower confidence, the distance is normalized, setting a threshold at 0.2. The quality of annotation is further enhanced through two subsequent filtering stages. In the first stage, lines that exhibit unusually high character rates, exceeding 37.5 Hz, are eliminated. The second stage encompasses another Whisper iteration; segments yielding a transcription of "Thank you." are excluded. These segments, which typically represent instrumental sections, are believed to originate from Whisper’s training on data similar to video transcripts.

Following the construction process outlined above, we ultimately obtained a multilingual lyric transcription dataset, MulJam, consisting of 6,031 songs with 182,429 lines and a total duration of 381.9 hours. The dataset’s statistical information and comparisons with existing ALT datasets are presented in Table 2.

Method	Jamendo	Hansen	DSing
TDNN-F [8]	76.37	77.59	19.60
CTDNN-SA [44]	66.96	78.53	14.96
Genre-informed AM [12]	50.64	39.00	56.90
MSTRE-Net [13]	34.94	36.78	15.38
DE2-segmented [45]	44.52	49.92	-
W2V2-ALT [22]	33.13	18.71	<b>12.99</b>
LyricWhiz (Ours)	<b>24.25</b>	<b>7.85</b>	13.78
w/o ChatGPT Ens.	<u>28.18</u>	<u>8.07</u>	15.22
w/o Whis. Prompt	33.21	8.75	<u>13.40</u>

**Table 3.** The WERs (%) of various ALT systems, including ablation methods, on multiple datasets. Note that W2V2-ALT is an in-domain baseline that natively train on DSing. The results of our method on Jamendo, Hansen are obtained from full-length transcription results, and the results on DSing are obtained from utterance-level segments.

To our best knowledge, MulJam is the first publicly available large-scale dataset for multilingual lyrics transcription without copyright restrictions. While DALI [7] is another large-scale music dataset featuring multilingual lyrics, its restricted access and strict licensing requirements limit its applicability for downstream tasks. In contrast, MulJam is free from copyright-related constraints and can be utilized without approval, as the audio can be legally downloaded directly from public sources without the need for approval, making it easily accessible. This even includes audio that is permitted for use in the development of commercial software. Researchers are permitted to legally modify our dataset for derivative works and redistribution, provided they cite our work and adhere to the CC BY-NC-SA license. Furthermore, in contrast to the imbalanced language distribution in DALI, where English songs account for over 80% of the total songs, our dataset includes a greater proportion of songs in other languages, which is advantageous for multilingual lyrics transcription.

#### 4. EXPERIMENTS

In this section, we first outline our experimental setup, including datasets and evaluation metrics. Next, we report lyrics transcription results on various benchmark datasets. We also conduct extensive ablation studies to verify the effectiveness of our methods. Finally, we demonstrate the reliability of our dataset through noise level estimation.

##### 4.1 Experimental Setup

**Datasets.** Our proposed method does not require any training; thus, we directly test it on several accessible lyric transcription benchmark datasets, including Jamendo [46], Hansen [15], MUSDB18 [17], DSing [8]. Among these, Jamendo, Hansen, and DSing are widely used test datasets in music transcription. MUSDB18, originally a dataset for music source separation, contains 150 rock-pop songs. The authors in [17] provided line-level lyric annotations for MUSDB18, making it a challenging real-world dataset for lyric transcription. Additionally, we manually collected 40 multilingual songs with lyrics annotations from MTG-

Method	a)	b)	c)
CTDNN-SA-mixture [17]	76.06	78.44	89.24
Ours-mixture	<b>50.90</b>	<b>47.04</b>	<b>50.70</b>
CTDNN-SA-vocals [17]	37.83	30.85	58.45
Ours-vocals	<b>26.29</b>	<b>25.27</b>	<b>33.30</b>

**Table 4.** The WERs (%) of our method and baseline [17] on three subsets of annotated MUSDB18. The results of our method are obtained from utterance-level segments.

Jamendo as a test set for the proposed dataset, which can be used to validate the reliability of our proposed dataset via transcription accuracy.

**Evaluation.** We report the Word Error Rate (WER) as the evaluation metric, which is the ratio of the total number of insertions, substitutions, and deletions with respect to the total number of words. We calculate the average WER on the test sets. Since Whisper possesses the capability for long-form transcription, we directly evaluate entire songs using Jamendo, Hansen, and the multilingual test set. We perform utterance-level evaluations on MUSDB18 and DSing since they only have utterance-level annotations. We discovered that many songs in these evaluation datasets are problematic, such as incorrect lyric annotations and excessively short song segments. One notable problem is that sometimes there are prominent harmony parts in the background of a song. However, it is not provided in the lyric annotations (e.g., Adele’s “Rolling in the Deep”). LyricWhiz is powerful enough to transcript both the leading vocal and the background vocal with high accuracy. Therefore, we removed these problematic data from our evaluations. Finally, we normalize the transcription results to match the standardized ground truths. We remove all special Unicode characters, such as emojis. All text is converted to lowercase, and numeric characters are converted to their alphabetic correspondence.

**Budget.** To ensure fast and multi-round inference of the Whisper-large model on various datasets, including the large-scale MTG-Jamendo dataset, we conducted our experiments concurrently on a server with 8xA100 80G GPUs. It takes approximately 9 hours to complete one round of inference, and each process uses up to 12G VRAM. The vocal probability threshold is set to 0.07 for PANNs-based vocal event detection. To carry out contextualized post-processing using ChatGPT, we invested a total of US\$2,000 on GPT-4 API for the entire project.

##### 4.2 Comparative Experiments

In order to verify the superiority of our approach, we compare it with several previous studies on benchmark datasets. W2V2-ALT [22], a transfer learning method based on ASR self-supervised models, represents the current state-of-the-art in lyric transcription tasks. In our experiments, we primarily compare our method with W2V2-ALT, as well as other previous methods. The experimental results, as shown in Table 3, indicate that our method achieves the best performance on Jamendo and Hansen and the second-best performance on DSing. In long-form



transcription datasets such as Jamendo and Hansen, our method significantly outperforms all previous approaches due to the strong contextual memory capabilities of both Whisper and ChatGPT. Furthermore, our method also leads by a considerable margin on MUSDB18, shown in Table 4, demonstrating the robust performance and resilience of our proposed method in more diverse and complex musical scenarios. It is worth noting that our method did not surpass previous results on the DSing dataset, which we attribute to two factors. First, previous models were trained on the DSing training set, making the DSing test set an in-distribution dataset for the models, while our approach does not require any training and directly employs large-scale ASR models for zero-shot lyric transcription. Second, the segmented evaluation on DSing results in the loss of contextual information, which consequently leads to inaccurate transcriptions.

### 4.3 Ablation Studies

To further substantiate the efficacy of each component within our proposed approach, we conducted comprehensive ablation experiments.

**Whisper Prompt.** In our experiments, we investigate the Whisper prompt mechanism and test various prompts. First, we construct a complex prompt following the format of ChatGPT prompts, including task descriptions, format specifications, and specific requirements. We then gradually reduce the constituent elements of the prompt and observe the results. We discover that, unlike general large language models, Whisper has weaker task understanding capabilities for complex prompts and can only comprehend shorter task prompts. In practice, using the simplest prompt “lyrics:” yielded the best results. For multilingual transcription, we translate “lyrics:” into the corresponding language. As shown in Table 3, the designed prompt performs better in long-form transcription scenarios, assisting the model in producing meaningful lyrics for difficult tasks. However, its performance is less effective at the utterance level, possibly because predicting a single line of lyrics does not require additional contextual information.

**ChatGPT Ensemble.** In order to confirm that ChatGPT can analyze and infer the most accurate version of lyrics, we first conduct a simple experiment. In this experiment, we add the ground truth lyrics to the predicted results and input them together into ChatGPT for ensembling. We then calculate the proportion of times ChatGPT ultimately chose the ground truth. If ChatGPT is able to choose the most accurate lyrics, *i.e.*, the ground truth, the final proportion should be close to 100%. The computed results on the Hansen dataset is 72.7% for ground truth data, which is sufficient to demonstrate that ChatGPT can make correct choices based on the constructed prompt and input lyrics. As further observed in Table 3, ChatGPT ensembling is particularly effective for long-form lyric transcription, suggesting that ChatGPT requires contextual information (the content of preceding and following lyrics, as well as the content of different versions of predicted lyrics) for inference. In contrast, utterance-level lyric inputs lack

Language	Songs <sub>train</sub>	Songs <sub>test</sub>	WER <sub>test</sub>
English	3,791	20	21.86
French	1,030	7	26.64
Spanish	620	5	22.54
Italian	311	3	44.01
Russian	147	4	39.18
German	132	1	25.43
Overall	6,031	40	26.26

**Table 5.** The distribution of our dataset and WERs (%) on test set. We manually constructed a test set of 40 songs following the language distribution of the collected training set. Then, we applied our proposed method to the test set and computed the WER.

both context and diversity among different prediction results, leading to inferior performance.

### 4.4 Dataset Analysis

In order to demonstrate the reliability of the dataset constructed using Whisper and ChatGPT on MTG-Jamendo, we manually create a multilingual test set for noise level estimation. Specifically, we first select six languages from the intersection of the languages in MTG and those in which Whisper performs best. We then conduct a stratified sampling of 40 songs on Jamendo and manually annotate their lyrics. We use these 40 songs as a test set, assessing the WER to estimate the noise level of our collected dataset. Table 5 presents the number of songs in each language and the WER results for the test set, where our method achieves decent WER levels for the majority of languages. As our goal is to construct a large-scale, multilingual dataset for weak supervision, our method’s transcription results are acceptable. Furthermore, we have not implemented specific normalization for multilingual transcription results, such as removing diacritical marks, which could be employed to enhance performance.

## 5. CONCLUSION

This paper presents LyricWhiz, a novel zero-shot automatic lyrics transcription system excelling in various datasets and music genres. Combining Whisper and GPT-4, our approach significantly reduces WER in English and efficiently transcribes multiple languages. LyricWhiz further generates the first publicly accessible, large-scale, multilingual lyrics dataset with a human-annotated subset for noise level estimation and evaluation. The successful integration of the large speech model and large language model in LyricWhiz offers a novel avenue for traditional Music Information Retrieval (MIR) tasks, as previous task-specific solutions are being eclipsed by general-purpose models. Notably, large language models have demonstrated their superior language understanding abilities across various tasks. Hence, we anticipate further applications of large language models to a broader spectrum of music-related domains, such as text-to-music generation, to enhance the performance of various models.

## 6. ACKNOWLEDGEMENTS

We gratefully acknowledge the dataset post-processing work described in Section 3.3 offered by Jiawen Huang. Jiahao Pan and Wei Xue were supported by the Theme-based Research Scheme (T45-205/21-N) and Early Career Scheme (ECS-HKUST22201322), Research Grants Council of Hong Kong. Yinghao Ma is a research student at the UKRI Centre for Doctoral Training in Artificial Intelligence and Music, supported by UK Research and Innovation [grant number EP/S022694/1]. Yizhi Li is fully funded by an industrial PhD studentship (Grant number: 171362) from the University of Sheffield, UK.

## 7. REFERENCES

- [1] A. Tsapras, “Lyrics-based music genre classification using a hierarchical attention network,” *arXiv preprint arXiv:1707.04678*, 2017.
- [2] H. Fujihara, M. Goto, and J. Ogata, “Hyperlinking Lyrics: A method for creating hyperlinks between phrases in song lyrics.” in *ISMIR*, 2008, pp. 281–286.
- [3] T. Hosoya, M. Suzuki, A. Ito, S. Makino, L. A. Smith, D. Bainbridge, and I. H. Witten, “Lyrics recognition from a singing voice based on finite state automaton for music information retrieval.” in *ISMIR*, 2005, pp. 532–535.
- [4] P. Knees and M. Schedl, “A survey of music similarity and recommendation from music context data,” *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, vol. 10, no. 1, pp. 1–21, 2013.
- [5] E. Çano and M. Morisio, “MoodyLyrics: A sentiment annotated lyrics dataset,” in *Proceedings of the 2017 international conference on intelligent systems, metaheuristics & swarm intelligence*, 2017, pp. 118–124.
- [6] P. Dhariwal, H. Jun, C. Payne, J. W. Kim, A. Radford, and I. Sutskever, “Jukebox: A generative model for music,” *arXiv preprint arXiv:2005.00341*, 2020.
- [7] G. Meseguer-Brocal, A. Cohen-Hadria, and G. Peeters, “DALI: a large dataset of synchronized audio, lyrics and notes, automatically created using teacher-student machine learning paradigm.” in *19th International Society for Music Information Retrieval Conference, ISMIR*, Ed., September 2018.
- [8] G. R. Dabike and J. Barker, “Automatic lyric transcription from karaoke vocal tracks: Resources and a baseline system.” in *Interspeech*, 2019, pp. 579–583.
- [9] P. Törnberg, “ChatGPT-4 outperforms experts and crowd workers in annotating political twitter messages with zero-shot learning,” *arXiv preprint arXiv:2304.06588*, 2023.
- [10] Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, and M. D. Plumbley, “PANNs: Large-scale pretrained audio neural networks for audio pattern recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2880–2894, 2020.
- [11] C. Gupta, H. Li, and Y. Wang, “Automatic pronunciation evaluation of singing.” in *Interspeech*, 2018, pp. 1507–1511.
- [12] C. Gupta, E. Yılmaz, and H. Li, “Automatic lyrics alignment and transcription in polyphonic music: Does background music help?” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 496–500.
- [13] E. Demirel, S. Ahlbäck, and S. Dixon, “MSTRE-Net: Multistreaming acoustic modeling for automatic lyrics transcription,” *arXiv preprint arXiv:2108.02625*, 2021.
- [14] G. Meseguer-Brocal, A. Cohen-Hadria, and G. Peeters, “Creating DALI, a large dataset of synchronized audio, lyrics, and notes,” *Transactions of the International Society for Music Information Retrieval*, vol. 3, no. 1, 2020.
- [15] J. K. Hansen and I. Fraunhofer, “Recognition of phonemes in a-cappella recordings using temporal patterns and mel frequency cepstral coefficients,” in *9th Sound and Music Computing Conference (SMC)*, 2012, pp. 494–499.
- [16] R. M. Bittner, K. Pasalo, J. J. Bosch, G. Meseguer-Brocal, and D. Rubinstein, “vocado: A dataset of solo vocals with  $f_0$ , note, and lyric annotations,” *arXiv preprint arXiv:2110.05580*, 2021.
- [17] K. Schulze-Forster, C. S. Doire, G. Richard, and R. Badeau, “Phoneme level lyrics alignment and text-informed singing voice separation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 2382–2395, 2021.
- [18] X. Gao, C. Gupta, and H. Li, “Genre-conditioned acoustic models for automatic lyrics transcription of polyphonic music,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 791–795.
- [19] S. Basak, S. Agarwal, S. Ganapathy, and N. Takahashi, “End-to-end lyrics recognition with voice to singing style transfer,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 266–270.
- [20] C. Zhang, J. Yu, L. Chang, X. Tan, J. Chen, T. Qin, and K. Zhang, “PDAugment: Data augmentation by pitch and duration adjustments for automatic lyrics transcription,” *arXiv preprint arXiv:2109.07940*, 2021.
- [21] A. M. Kruspe and I. Fraunhofer, “Training phoneme models for singing with “songified” speech data.” in *ISMIR*, 2015, pp. 336–342.

- [22] L. Ou, X. Gu, and Y. Wang, “Transfer learning of wav2vec 2.0 for automatic lyric transcription,” in *ISMIR*, 2022.
- [23] R. Tang, K. Kumar, G. Yang, A. Pandey, Y. Mao, V. Belyaev, M. Emmadi, C. Murray, F. Ture, and J. Lin, “SpeechNet: Weakly supervised, end-to-end speech recognition at industrial scale,” *arXiv preprint arXiv:2211.11740*, 2022.
- [24] S. Schneider, A. Baevski, R. Collobert, and M. Auli, “wav2vec: Unsupervised pre-training for speech recognition,” *arXiv preprint arXiv:1904.05862*, 2019.
- [25] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [26] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, “HuBERT: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [27] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, “WavLM: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [28] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” *arXiv preprint arXiv:2212.04356*, 2022.
- [29] C. Wang, S. Chen, Y. Wu, Z. Zhang, L. Zhou, S. Liu, Z. Chen, Y. Liu, H. Wang, J. Li *et al.*, “Neural codec language models are zero-shot text to speech synthesizers,” *arXiv preprint arXiv:2301.02111*, 2023.
- [30] A. Agostinelli, T. I. Denk, Z. Borsos, J. Engel, M. Verzetti, A. Caillon, Q. Huang, A. Jansen, A. Roberts, M. Tagliasacchi *et al.*, “MusicLM: Generating music from text,” *arXiv preprint arXiv:2301.11325*, 2023.
- [31] M. C. McCallum, F. Korzeniowski, S. Oramas, F. Gouyon, and A. F. Ehmann, “Supervised and unsupervised learning of audio representations for music understanding,” *arXiv preprint arXiv:2210.03799*, 2022.
- [32] C. Donahue, A. Caillon, A. Roberts, E. Manilow, P. Esling, A. Agostinelli, M. Verzetti, I. Simon, O. Pietquin, N. Zeghidour *et al.*, “SingSong: Generating musical accompaniments from singing,” *arXiv preprint arXiv:2301.12662*, 2023.
- [33] Y. Li, R. Yuan, G. Zhang, Y. Ma, C. Lin, X. Chen, A. Ragni, H. Yin, Z. Hu, H. He *et al.*, “MAP-Music2Vec: A simple and effective baseline for self-supervised music audio representation learning,” *arXiv preprint arXiv:2212.02508*, 2022.
- [34] Y. Li, R. Yuan, G. Zhang, Y. Ma, X. Chen, H. Yin, C. Lin, A. Ragni, E. Benetos, N. Gyenge *et al.*, “MERT: Acoustic music understanding model with large-scale self-supervised training,” *arXiv preprint arXiv:2306.00107*, 2023.
- [35] Y. Shen, K. Song, X. Tan, D. Li, W. Lu, and Y. Zhuang, “HuggingGPT: Solving ai tasks with chatgpt and its friends in huggingface,” *arXiv preprint arXiv:2303.17580*, 2023.
- [36] T. Schick, J. Dwivedi-Yu, R. Dessì, R. Raileanu, M. Lomeli, L. Zettlemoyer, N. Cancedda, and T. Scialom, “Toolformer: Language models can teach themselves to use tools,” *arXiv preprint arXiv:2302.04761*, 2023.
- [37] P. Liu, W. Yuan, J. Fu, Z. Jiang, H. Hayashi, and G. Neubig, “Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing,” *ACM Computing Surveys*, vol. 55, no. 9, pp. 1–35, 2023.
- [38] J. White, S. Hays, Q. Fu, J. Spencer-Smith, and D. C. Schmidt, “ChatGPT prompt patterns for improving code quality, refactoring, requirements elicitation, and software design,” *arXiv preprint arXiv:2303.07839*, 2023.
- [39] J. White, Q. Fu, S. Hays, M. Sandborn, C. Olea, H. Gilbert, A. Elnashar, J. Spencer-Smith, and D. C. Schmidt, “A prompt pattern catalog to enhance prompt engineering with ChatGPT,” *arXiv preprint arXiv:2302.11382*, 2023.
- [40] T. Shin, Y. Razeghi, R. L. Logan IV, E. Wallace, and S. Singh, “AutoPrompt: Eliciting knowledge from language models with automatically generated prompts,” *arXiv preprint arXiv:2010.15980*, 2020.
- [41] K. Shum, S. Diao, and T. Zhang, “Automatic prompt augmentation and selection with chain-of-thought from labeled data,” *arXiv preprint arXiv:2302.12822*, 2023.
- [42] J. Wei, X. Wang, D. Schuurmans, M. Bosma, E. Chi, Q. Le, and D. Zhou, “Chain of thought prompting elicits reasoning in large language models,” *arXiv preprint arXiv:2201.11903*, 2022.
- [43] V. I. Levenshtein *et al.*, “Binary codes capable of correcting deletions, insertions, and reversals,” in *Soviet physics doklady*, vol. 10, no. 8. Soviet Union, 1966, pp. 707–710.



- [44] E. Demirel, S. Ahlbäck, and S. Dixon, “Automatic lyrics transcription using dilated convolutional neural networks with self-attention,” in *2020 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2020, pp. 1–8.
- [45] E. Demirel, S. Ahlbäck, and S. Dixon, “Low resource audio-to-lyrics alignment from polyphonic music recordings,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 586–590.
- [46] D. Stoller, S. Durand, and S. Ewert, “End-to-end lyrics alignment for polyphonic music using an audio-to-character recognition model,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 181–185.