

# FRUITSMUSIC: A REAL-WORLD CORPUS OF JAPANESE IDOL-GROUP SONGS

Hitoshi Suda<sup>1</sup> Shunsuke Yoshida<sup>2</sup> Tomohiko Nakamura<sup>1</sup> Satoru Fukayama<sup>1</sup> Jun Ogata<sup>1</sup>

<sup>1</sup> National Institute of Advanced Industrial Science and Technology (AIST), Tokyo, Japan

<sup>2</sup> The University of Tokyo, Tokyo, Japan

{suda.h, shunsuke.yoshida, tomohiko-nakamura, s.fukayama, jun.ogata}@aist.go.jp

## ABSTRACT

This study presents FruitsMusic, a metadata corpus of Japanese idol-group songs in the real world, precisely annotated with who sings what and when. Japanese idol-group songs, vital to Japanese pop culture, feature a unique vocal arrangement style, where songs are divided into several segments, and a specific individual or multiple singers are assigned to each segment. To enhance singer diarization methods for recognizing such structures, we constructed FruitsMusic as a resource using 40 music videos of Japanese idol groups from YouTube. The corpus includes detailed annotations, covering songs across various genres, division and assignment styles, and groups ranging from 4 to 9 members. FruitsMusic also facilitates the development of various music information retrieval techniques, such as lyrics transcription and singer identification, benefiting not only Japanese idol-group songs but also a wide range of songs featuring single or multiple singers from various cultures. This paper offers a comprehensive overview of FruitsMusic, including its creation methodology and unique characteristics compared to conversational speech. Additionally, this paper evaluates the efficacy of current methods for singer embedding extraction and diarization in challenging real-world conditions using FruitsMusic. Furthermore, this paper examines potential improvements in automatic diarization performance through evaluating human performance.

## 1. INTRODUCTION

In Japanese pop culture, an *idol* is a performer who engages in dancing, singing, and entertaining fans [1]. In the culture, idols frequently participate in activities, such as concerts and television programs, as members of idol groups. One of the most renowned contemporary idol groups is AKB48, which has 40 single compact discs (CDs) that are million-sellers, as certified by The Recording Industry Association of Japan<sup>1</sup>. FRUITS ZIPPER has

<sup>1</sup> [https://www.riaj.or.jp/f/data/cert/gd\\_search.html](https://www.riaj.or.jp/f/data/cert/gd_search.html)



emerged as another notable group comprising seven girls and being awarded the Best New Artist at the Japan Record Awards 2023, the most prestigious accolade in Japanese music culture [2]. Not only can fans attend concerts, but they can also interact with the idols at handshaking events (*Akushukai*) or bonus events (*Tokutenkai*), where the fans can forge deep connections with the idols [3].

Idol-group songs feature several unique characteristics. One notable characteristic is *song division*, also called *utawari* in Japanese [4, 5]. This approach involves a dynamic vocal arrangement where the singing roles shift throughout the song; individual members may take turns singing solo lines, or multiple members may sing together in unison. In particular, the entire group often sings together in the chorus sections, known as *sabi*. Song division is chosen intentionally to maximize the charm and attractiveness of each idol and song. Therefore, the analysis of song division is crucial for understanding the structure and expression of songs, as well as the creators’ intentions.

Song division plays a crucial role also in shaping audience participation through chants and shouts, known as *calls* and *mixes*, which are indispensable elements of idol-group concerts [6]. Fans spontaneously create these chants and shouts, reflecting the song’s structure, musical intensity, and song division, specifically which member is assigned to sing at any given moment. Furthermore, song division significantly influences music videos and concert recordings produced by idol groups, demonstrating its pivotal role in producing and appreciating idol music content.

As previously described, song division is crucial for understanding and enjoying the musical compositions of idol groups. To aid fans’ comprehension, some idol groups release official charts showing how songs are divided among members. For Korean pop groups with similar features to Japanese idol groups, several fans create *line distribution* videos. These videos, widely viewed on platforms like YouTube and TikTok, visualize the structures of song division, facilitating a deeper understanding. Therefore, developing techniques for recognizing song division will help fans enjoy the music compositions and enhance their interaction with idols. In addition, such advancements will support creators in promoting idol groups.

The task to estimate song division, i.e., *who sings when*, within a music signal is known as *singer diarization*. This technique has been inspired by speaker diarization, which identifies *who speaks when* in conversational speech [7–9]. The singer diarization technique was initially introduced

to analyze folk music and has been adapted for Japanese idol-group songs [4]. However, existing research has not examined songs from real-world idol groups but from idol-themed games and anime. These game and anime songs generally belong to narrower genres, feature simpler song division structures, and have vocals that are easier to distinguish, thanks to the distinctive voice qualities of the voice actors. Further research indicates that in-the-wild audio signals can improve diarization performance in real-world settings, even with small datasets [4, 10]. Consequently, compiling a dataset featuring songs from real-world idol groups is critical for developing practical applications targeting pop culture.

This study addresses the demand for a practical dataset in music information retrieval (MIR) by constructing a new corpus, FruitsMusic. This corpus consists of detailed annotations about *who sings what and when* in real-world songs performed by Japanese idol groups from YouTube, enabling the advancement of singer diarization techniques and their assessments. Beyond singer diarization, FruitsMusic also advances various MIR techniques such as lyrics transcription [11, 12], emotion classification [13, 14], singer identification [15, 16], and singer-based music search [17], for not only Japanese idol-group songs but also a wide array of musical pieces featuring single or multiple singers from different cultures. A significant advantage of FruitsMusic is its focus on real idol groups, allowing for evaluations in challenging scenarios and enhancing the applicability of MIR techniques in the real world. This paper details the structure, development methodology, and unique characteristics of FruitsMusic. The paper also demonstrates the applications of evaluating existing methods in two MIR tasks, singer embedding extraction and diarization, in real-world scenarios.

## 2. STRUCTURE AND CONSTRUCTION METHODOLOGY OF FRUITSMUSIC

In this study, we constructed FruitsMusic (Corpus of Fully Real-World Popular Idol-group Songs from YouTube Videos for Music Information Processing) aimed at developing and evaluating various MIR techniques. This corpus is a collection of annotations for 163 minutes of music video content on YouTube, detailing *who sings what and when*. The corpus includes annotations for 40 songs performed by 18 different groups, featuring a total of 122 unique female singers, all approximately 20 years of age. The corpus is available at <https://huggingface.co/datasets/fruits-music/fruits-music><sup>2</sup>.

### 2.1 Related Works

Several corpora derived from YouTube have been constructed across various research fields. The key advantage of this approach is the utilization of a wide range of real-world video and audio content.

For example, ActivityNet and YouTube-8M are benchmark datasets widely used in video processing [18, 19].

<sup>2</sup> This paper has been written based on FruitsMusic version 1.2.0.

```
{
  "id": "XXm01",
  "youtubeId": "YouTube ID",
  "type": "music_video",
  "singerIds": ["XXs01", "XXs02", "XXs04", "XXs05", "XXs06"],
  "title": "Song Title",
  "songStartsAt": 0,
  "duration": 216128,
  "states": [
    {
      "start": 1869,
      "end": 17233,
      "singers": [0, 1, 2, 3, 4],
      "lyrics": "Lyrics 1",
      "realLyrics": null
    },
    {
      "start": 22543,
      "end": 26930,
      "singers": [1],
      "lyrics": "Lyrics 2",
      "realLyrics": null
    }
  ]
}
```

(a) JSON file

```
SPEAKER XXm01 1 1.869 15.364 <NA> <NA> XXs01 <NA> <NA>
SPEAKER XXm01 1 1.869 15.364 <NA> <NA> XXs02 <NA> <NA>
SPEAKER XXm01 1 1.869 15.364 <NA> <NA> XXs04 <NA> <NA>
SPEAKER XXm01 1 1.869 15.364 <NA> <NA> XXs05 <NA> <NA>
SPEAKER XXm01 1 1.869 15.364 <NA> <NA> XXs06 <NA> <NA>
SPEAKER XXm01 1 22.543 4.387 <NA> <NA> XXs02 <NA> <NA>
```

(b) RTTM file

**Table 1.** An example of JSON and RTTM files.

Similarly, YouTube-ASL, a large-scale American Sign Language corpus, originates from YouTube [20].

In the field of audio processing, several corpora have utilized YouTube videos. AudioSet, for instance, is widely adopted for recognizing and detecting audio events [21]. VoxLingua107 covers 6,628 hours of speech across 107 languages and is helpful to language detection [22]. Further, JTubeSpeech consists of extensive Japanese speech data from YouTube and helps the development of diverse speech processing techniques [23]. Similarly, YODAS consists of 500,000 hours of speech in over 100 languages and makes multilingual speech processing techniques applicable in the wild [23, 24]. Coco-Nut is another corpus with subjective descriptions of voices, designed for controlling speaker identity based on text prompts [25].

These prior works underscore the effectiveness of YouTube-based corpora, which we also adopted in this study. Our corpus focuses especially on accuracy and reliability, which are less emphasized in these prior corpora. In addition, the video-based nature of FruitsMusic facilitates multimodal processing, such as multimodal diarization [26]. Note that these prior corpora have been curated to protect individual privacy rights by excluding personal information, and FruitsMusic also maintains these ethical standards.

### 2.2 Structure of the Corpus

FruitsMusic includes annotations in JavaScript Object Notation (JSON) format, Rich Transcription Time Marked (RTTM) files for diarization, and text files of lyrics. Table 1 presents an example of JSON and RTTM files.

### 2.2.1 JSON Files

The JSON files include the following information:

- **Song ID.** This field is formed by combining a two-character idol-group ID, the letter “m”, and a two-digit ID.
- **Video ID on YouTube.**
- **Type of the video.** This field is either of `music_video`, `middle_music_video`, or `dance_practice`. The names of these types are derived from traditions in Japanese idol culture.
- **List of singer IDs.** Each ID is formed by combining a two-character idol-group ID, the letter “s”, and a two-digit ID.
- **Song title.** This field aims at natural language processing (NLP) tasks.
- **Start time and duration of the song.** The videos may contain content beyond songs, such as comments from idols. This information is provided to help filter out such content.
- **Singing states.** This is a list of the start and end times of the segment, the singers assigned to the segment, and the lyrics. The `lyrics` field contains the official lyrics, which may differ from the actual lyrics sung. In such cases, the `realLyrics` field is used.

The time and duration fields are annotated in milliseconds.

### 2.2.2 RTTM Files

The RTTM format is specially designed for speaker diarization tasks, identifying *who speaks when* [27]. Table 1b presents an example of an RTTM file. Within this format, each line details the start time and duration of the segment, as well as the singer’s ID. For simultaneous singing, the format allocates a separate line to each singer, resulting in multiple lines corresponding to the number of singers.

### 2.2.3 Text Files of Lyrics

Lyrics lines may be duplicated in the JSON files to precisely represent *who sings what and when* (e.g., DRm03). As a result, extracting lyrics from JSON files is not straightforward. To support the development and assessment of techniques involving lyrics, such as lyrics transcription, FruitsMusic provides separate text files of lyrics.

## 2.3 Subsets

FruitsMusic is split into Subset A and Subset B. Subset A is designed mainly for training, and Subset B is for evaluation. However, both subsets can be arbitrarily used for various purposes. Subset A contains 32 songs, while Subset B has 8 songs. To ensure unbiased evaluation, Subset B does not contain any singers from Subset A, and each group in Subset B contributes only one song. The songs in Subset B were chosen to cover various genres (dance, rock, synthpop, etc.) and division styles. Also, groups in Subset B are generally less famous than those in Subset A, which helps ensure fairer and less biased human evaluation.

|                           | CHiME-5  | FruitsMusic |
|---------------------------|----------|-------------|
| Average audio length      | 9031 s   | 244 s       |
| # Speakers                | 4        | 4-9         |
| Average segment length    | 2.11 s   | 4.44 s      |
| Total length per speaker  | 1159.6 s | 15.9 s      |
| Segments without speakers | 22.3%    | 23.9%       |
| Solo segments             | 51.4%    | 42.6%       |
| Multiple-speaker segments | 26.4%    | 33.5%       |
| Segments with 3+ speakers | 6.4%     | 26.5%       |

**Table 2.** Comparison of FruitsMusic with the CHiME-5 dataset [28], a conversational speech dataset. The “Total length per speaker” row indicates the average total duration per speaker in each audio.

## 2.4 Song Selection

We meticulously selected the songs for FruitsMusic to ensure the corpus’s reliability and usefulness. To achieve accurate annotations, we initially gathered extensive knowledge about the idol groups. We then used reliable sources, including concert recordings and official announcements, for information. Additionally, to support applications like singer diarization, each singer has at least one solo section within the database. Moreover, we assign each singer to only one group in FruitsMusic. While idols may participate in multiple groups or move between groups in reality, we avoid such complexities in this database. FruitsMusic focuses solely on contemporary songs released from 2022 onwards to reflect the latest music trends.

## 2.5 Rules

This corpus has been constructed using copyrighted materials. Users are required to follow the licensing agreement specified in the corpus documentation to protect the rights of creators and idols. The agreement sets three major rules. First, the copyrighted content of this corpus, such as lyrics texts, is not intended for appreciation or entertainment. Second, the corpus cannot be used to develop or enhance generative artificial intelligence (AI) techniques, such as singing voice synthesis, voice conversion, lyrics generation, and music creation. However, users can utilize the corpus for recognition or information extraction tasks, including lyrics recognition, singer embedding extraction, and assessing the naturalness of lyrics or music. Third, when citing this corpus in any media, including academic works and presentations, users are required to identify both the groups and the singers using the provided IDs and refrain from using their real names. If the mention of song names is not essential for the discussion, users are also required to refer to them by their respective IDs.

## 3. COMPARISON WITH CONVERSATIONAL DATASET

This section compares FruitsMusic with the CHiME-5 dataset, a conversational speech dataset designed for The

5th CHiME Speech Separation and Recognition Challenge [28], to explore the differences between conversational speech and songs with song division. The CHiME-5 dataset contains 20 conversational speech instances, each from four speakers. Table 2 shows the comparison results, considering all subsets of both CHiME-5 and FruitsMusic.

Initially, the average audio length in FruitsMusic is significantly shorter than in CHiME-5. Unlike conversational speech, which is not limited by specific length constraints, the duration of songs is tightly controlled by the structure of the musical compositions. Furthermore, the average total duration of speech segments per speaker in FruitsMusic is extremely shorter than in CHiME-5. This difference arises from the shorter overall audio length and the larger number of singers in FruitsMusic. Since solo segments play a key role in capturing singer characteristics, developing singer identification techniques under these challenging conditions, different from conversational speech, is essential for improving singer diarization and other MIR systems for songs featuring multiple singers.

The comparison reveals a noteworthy difference in the frequency of simultaneous speakers between CHiME-5 and FruitsMusic. In particular, sections featuring 3 or more singers in FruitsMusic are significantly longer than in CHiME-5. This indicates that the methods that treat overlapping speech as segments with two speakers, often adopted in speaker diarization [29], cannot be directly applied to singer diarization. Furthermore, about 60% of segments with multiple singers feature vocals from only a subset of the entire group. Hence, the assumption that all singers are present in overlapped segments proves ineffective for singer diarization; it is crucial to accurately and independently determine the vocal activity of each singer.

#### 4. APPLICATION 1: SINGER EMBEDDINGS

Singer embeddings are multidimensional vectors that capture each singer’s unique vocal traits. In singing information processing, high-quality singer embeddings are crucial for enhancing the performance of tasks involving singers, such as singer identification, voice matching, and singer diarization. This section evaluates two types of embeddings extracted from song segments by a specific group and discusses the effectiveness of each extraction technique in real-world scenarios. This section visualizes these embeddings to understand their effectiveness in distinguishing singers and provides a numerical analysis of the clustering performance based on singers.

In our evaluation, we compare two types of singer embeddings. The first type involves x-vectors, traditional yet effective speaker embeddings derived from deep neural networks (DNNs) for speaker identification [30]. Specifically, we utilize an x-vector extractor `microsoft/wavlm-base-plus-sv`<sup>3</sup>, which incorporates WavLM, a large-scale pre-trained model based on self-supervised learning [31]. Second, we evaluated embeddings based on ECAPA-TDNN, an enhanced time de-

lay neural network (TDNN) in x-vector extractors [32]. ECAPA-TDNN-based embeddings have been proven to show remarkable performance in speaker recognition and diarization [32–34]. We used an ECAPA-TDNN model provided by SeechBrain<sup>4</sup> [35]. In addition, this evaluation considers both mixed and vocal signals, with the latter extracted using Demucs, an open-source music source separation tool [36, 37]. We utilized the `htdemucs_ft` model, a fine-tuned version of the Hybrid Transformer Demucs, renowned for its state-of-the-art performance in music source separation.

We focus on the group KF, which comprises seven members and has eight songs, the most available on FruitsMusic. For this study, we selected segments where a singer performs solo for over 2 seconds. On average, each singer has 20 segments, totaling approximately 101 seconds of solo performance.

As objective evaluation metrics,  $F$  values are calculated to benchmark clustering efficacy [38]. Here, the  $F$  value is the harmonic mean of two metrics: purity  $P$  and inverse purity  $I$ . The  $P$  and  $I$  are defined as follows:

$$P = \frac{1}{N} \sum_i \max_j |C_i \cap S_j|, \text{ and} \quad (1)$$

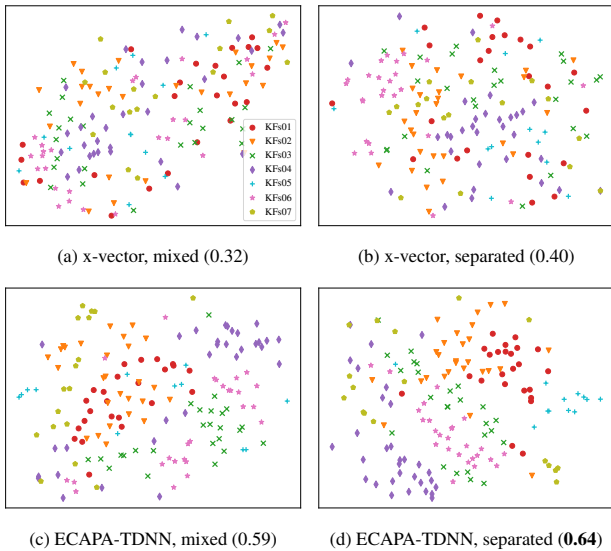
$$I = \frac{1}{N} \sum_j \max_i |C_i \cap S_j|. \quad (2)$$

In these equations,  $C_i$  is the  $i$ -th cluster, and  $S_j$  is the set of the  $j$ -th singer’s samples. High  $F$  values indicate superior performance, with a theoretical maximum of 1. For this evaluation, spectral clustering [39] was performed to create 7 clusters, matching the number of singers. All the embeddings were  $L_2$ -normalized in advance.

Figure 1 shows the visualizations of the acquired embeddings by reducing their dimensions into two using  $t$ -SNE. The effectiveness of Demucs is confirmed across both extraction methods. Compared to x-vectors, embeddings derived with ECAPA-TDNN provide more expressive singer representations. Specifically, Figure 1d reveals that samples from certain singers, specifically KFs01 (represented in red circles) and KFs06 (in pink stars), tend to gather by singer identity. Hence, ECAPA-TDNN-based embeddings are proved to effectively capture singers’ unique identities even from short singing segments. This shows the advantages of the ECAPA-TDNN methodology over conventional TDNN in x-vector extractors. However, none of the plots show distinct clusters visibly forming, and the highest  $F$  value was only 0.64. This indicates that tasks like singer diarization and number estimation remain challenging using any embedding extractor evaluated. Since the ECAPA-TDNN model is trained with speech datasets, fine-tuning it with singing voice datasets will enhance its performance. Note that the separated vocal signals are distorted; therefore, using datasets with both clean and mixed or separated signals from real-world conditions will be effective.

<sup>3</sup> <https://huggingface.co/microsoft/wavlm-base-plus-sv>

<sup>4</sup> <https://huggingface.co/speechbrain/spkrec-ecapa-voxceleb>



**Figure 1.** *t*-SNE visualizations of singer embeddings from the idol group KF’s songs, where each color and shape represents a different singer. Captions detail the extraction methods and whether Demucs was applied. Values in parentheses represent  $F$  values, measuring the clustering performance.

## 5. APPLICATION 2: SINGER DIARIZATION

To evaluate the efficacy of FruitsMusic in training singer diarization models, we trained several models with Subset A of FruitsMusic and assessed their performance using Subset B. In this comparison, the number of singers was not given to the systems. Furthermore, we engaged a human evaluator to perform the manual diarization of songs in Subset B and discuss the potential advancements in automatic diarization performance.

### 5.1 Construction of a Synthesized Dataset

To improve the diarization performance, we utilized songs from commercial CDs in addition to FruitsMusic. This dataset consists of 272 songs performed by multiple singers, with a separate recording for each song and singer combination. For example, if three singers perform song A, each of the three singers has solo recordings: one by singer 1, another by singer 2, and a third by singer 3. On average, each song features 4.1 singers, resulting in a total of 1126 recordings. All the songs were sourced from idol-themed games and anime, and the singers were 129 unique female voice actresses. We executed source separation on all 1126 recordings using Demucs to generate isolated vocal and accompaniment signals.

We generated five song division patterns for each song, capping the number of singers to a maximum of seven. We applied voice activity detection (VAD) first and randomly assigned singers to each segment. During the assignment, a single singer was allocated to 60% of the segments, all singers to 23%, and random singers to the remaining 17%. We mixed the vocal signals based on the generated song division and combined the sepa-

rated accompaniment with the mixed vocal tracks to create the final mixture. In these generated songs with song division, singers perform in unison during segments with multiple singers. The VAD process used `pyannote/voice-activity-detection`<sup>5</sup>.

### 5.2 Evaluated Systems

In this experiment, the following systems are compared.

#### 5.2.1 SA-EEND with EDA

The first approach adopted Self-Attentive End-to-End Neural Diarization (SA-EEND) [10]. Since the number of singers for the evaluated signals was unknown, we used enhanced SA-EEND with Encoder-Decoder-based Attractors (EDA) [40]. The hyperparameters matched those of the CALLHOME dataset, as specified in the original publication [40]. The input signals were downsampled to 8000 Hz and were converted to monaural signals.

#### 5.2.2 `pyannote.audio`

The second method used `pyannote.audio`<sup>6</sup>, an open-source toolkit for speech processing tasks [29, 41]. The diarization workflow is structured as a pipeline process, incorporating PyanNet-based modules. To conduct this experiment, we fine-tuned the publicly available pre-trained model `pyannote/speaker-diarization-3.1`<sup>7</sup> using the prepared song datasets. This fine-tuning process adapted the segmentation models and optimized the thresholds for both segmentation and clustering. The input signals were downsampled to 16 000 Hz and converted to monaural signals.

#### 5.2.3 Human Evaluator

In addition to the automatic diarization approaches, we also engaged a human evaluator to perform manual singer diarization to gauge the achievable performance. The individual understands Japanese and often listens to Japanese pop music (about 60 hours a month), yet was completely unfamiliar with any of the songs in Subset B of FruitsMusic. To maintain the experiment’s integrity, we presented only the audio signals of the songs without any corresponding videos. The participant was allowed to use any external tool to aid in the diarization process but was explicitly restricted from searching for the songs on the internet.

### 5.3 Experimental Setup

As a training dataset, Subset A from FruitsMusic was used. The songs DRm01, Kfm01, Rgm01, SBm01, and SYm01 were designated for validation. The remaining songs, excluding three songs featuring nine singers, were allocated for training. Due to the highly extended training time required for SA-EEND with EDA for songs featuring more than seven singers, three songs with nine singers, VYm02,

<sup>5</sup> <https://huggingface.co/pyannote/voice-activity-detection>

<sup>6</sup> <https://github.com/pyannote/pyannote-audio>

<sup>7</sup> <https://huggingface.co/pyannote/speaker-diarization-3.1>



| System                    | Mixed        | Separated |
|---------------------------|--------------|-----------|
| <b>SA-EEND with EDA</b>   |              |           |
| Synthesized only          | 99.5%        | 101.3%    |
| Synthesized + FruitsMusic | 103.2%       | 83.8%     |
| <b>pyannotate.audio</b>   |              |           |
| Synthesized only          | 92.9%        | 69.9%     |
| Synthesized + FruitsMusic | 91.3%        | 50.3%     |
| Human                     | <b>22.7%</b> | —         |

**Table 3.** DER for Subset B in FruitsMusic with the several diarization systems.

VYm03, and XSm02, were excluded from the dataset. The loudness of all songs was normalized to  $-14$  LUFS.

The evaluation metric used was the diarization error rate (DER) [27], defined as:

$$\text{DER} = \frac{\sum_{s=1}^S d_s \left[ \max \left( N_s^{(\text{ref})}, N_s^{(\text{hyp})} \right) - N_s^{(\text{correct})} \right]}{\sum_{s=1}^S d_s N_s^{(\text{ref})}}. \quad (3)$$

Here,  $S$  is the total number of segments,  $d_s$  represents the duration of the  $s$ -th segment, and  $N_s^{(\text{ref})}$ ,  $N_s^{(\text{hyp})}$ , and  $N_s^{(\text{correct})}$  correspond to the number of ground-truth singers, estimated singers, and accurately identified singers in the  $s$ -th segment, respectively. According to this definition, DER can exceed 100%. The calculation of DER was performed with `dscore`<sup>8</sup>, an open-source tool. Due to the implementation of `dscore`, self-overlapped segments, which contain multiple recordings of the same singer, were normalized in the calculation process. The collar size, the time ignored in DER calculation around segment boundaries, was set to zero.

The model selection criterion was achieving the minimum DER on the validation set. For each condition, we developed two versions of the system: one trained on mixed signals and another trained on extracted vocal signals. The vocal signal extraction was performed using the `htdemucs_ft` model of `Demucs` [36, 37].

## 5.4 Results

Table 3 shows the DER of all the systems. The performance of the mixed signal systems is significantly inferior to that of the separated signal systems. In other words, across evaluated systems, `Demucs` effectively improved diarization performance; hence, a pipeline system combining source separation and diarization proved more effective than using a single system on mixed signals in the case of this evaluation. In both approaches, `SA-EEND` and `pyannotate.audio`, training with `FruitsMusic` significantly improved the overall performance, particularly for the separated signal systems. The results suggest that `FruitsMusic`, despite its smaller size, can significantly enhance diarization performance rather than relying solely on large-scale synthesized datasets.

<sup>8</sup> <https://github.com/nryant/dscore>

| System                  | BD | BI | JA       | JY       | MG       | QD       | SL       | TJ       |
|-------------------------|----|----|----------|----------|----------|----------|----------|----------|
| <b>SA-EEND with EDA</b> |    |    |          |          |          |          |          |          |
| w/o FruitsMusic         | 1  | 3  | 2        | 2        | 4        | 0        | <b>6</b> | 2        |
| w/ FruitsMusic          | 2  | 2  | 2        | 2        | 2        | 2        | 2        | 2        |
| <b>pyannotate.audio</b> |    |    |          |          |          |          |          |          |
| w/o FruitsMusic         | 3  | 5  | 3        | 3        | 3        | 3        | 3        | 3        |
| w/ FruitsMusic          | 7  | 7  | <b>7</b> | 7        | <b>7</b> | 6        | 7        | 7        |
| Human                   | 8  | 6  | 6        | <b>5</b> | <b>7</b> | <b>5</b> | <b>6</b> | <b>4</b> |
| Ground truth            | 9  | 4  | 7        | 5        | 7        | 5        | 6        | 4        |

**Table 4.** Estimated total number of singers derived from diarization results. All the systems used separated vocal signals using `Demucs`. Each column shows a song in Subset B. The suffixes “m01” of song IDs are omitted.

Table 4 shows the estimated number of singers included in the diarization results. The `SA-EEND`-based systems struggled to distinguish singers accurately. This seems due to the difficulties of naive DNN-based methods in distinguishing singer identities, as discussed in Section 4. On the other hand, `pyannotate.audio` demonstrated an almost invariant estimation of the number of singers. This indicates a potential overfitting to the training datasets, with the most common number of singers in the training set tending to dominate the predictions.

Among the evaluated systems, human performance was remarkably superior to the automatic diarization systems in terms of DER. Notably, a human evaluator accurately estimated the number of singers in 5 out of 8 songs. This demonstrates that humans can effectively distinguish individual singers’ voices even within mixed music signals. Therefore, these results proved a significant potential for improving both automatic singer identification and diarization performance.

## 6. CONCLUSION

This paper presents `FruitsMusic`, a novel corpus of precise annotations on *who sings what and when* in Japanese idol-group songs. The song selection and subset creation were meticulously conducted to facilitate unbiased evaluation and ensure usefulness across a wide range of genres, song division styles, and idol groups. The corpus can be applied to various MIR tasks, such as singer diarization, singer identification, and lyrics transcription. This paper showcases its applications in evaluating singer embedding extraction and diarization techniques. The results showed that distinguishing singers from short singing segments remains challenging, despite effective methods in speech processing. The paper also suggests potential advancements in automatic diarization performance by assessing human performance. We acknowledge significant existing areas for performance improvement in diverse MIR tasks, and we are confident that `FruitsMusic` has the potential to advance various techniques among them.

## 7. ACKNOWLEDGMENTS

This work was supported by JSPS KAKENHI Grant Number JP23K20017. This paper is based on results obtained from a project, JPNP20006, commissioned by the New Energy and Industrial Technology Development Organization (NEDO). This study was supported by the BRIDGE program of the Cabinet Office, Government of Japan.

## 8. ETHICS STATEMENT

The FruitsMusic corpus is derived from YouTube videos of Japanese idol groups and consists of annotations that detail *who sings what and when*. Hence, FruitsMusic has the potential to help develop singing voice synthesis (SVS) or voice conversion (VC) systems that replicate the voices of actual idols. Moreover, analysis of the videos within FruitsMusic enables associating the real names of groups and singers with their IDs. These scenarios raise potential concerns about infringing on the personality rights of the idols. Additionally, using FruitsMusic to construct or trigger other generative AI techniques, such as lyrics and music generation, may violate the rights of composers and creators. In light of these considerations, as described in Section 2.5, FruitsMusic enforces the stringent licensing agreement and requires all users to adhere to it when downloading and using the corpus.

## 9. REFERENCES

- [1] P. W. Galbraith and J. G. Karlin, "Introduction: The mirror of idols and celebrity," in *Idols and Celebrity in Japanese Media Culture*, P. W. Galbraith and J. G. Karlin, Eds., Aug. 2012, pp. 1–32.
- [2] "FRUITS ZIPPER was awarded the Best New Artist at the Japan Record Awards 2023 (in Japanese)," *Sankei Shimbun*, Dec. 2023.
- [3] S. Nozawa, "The concentration booth and the handshaking lane: Ideologies of the phatic," *International Journal of the Sociology of Language*, vol. 2023, no. 284, pp. 15–36, Nov. 2023.
- [4] H. Suda, D. Saito, S. Fukayama, T. Nakano, and M. Goto, "Singer diarization for polyphonic music with unison singing," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 1531–1545, May 2022.
- [5] Y. Okada, Ed., *Living with Idols (in Japanese)*. Pot Publishing, Jul. 2013.
- [6] W. Xie, "Japanese "idols" in trans-cultural reception: the case of AKB48," in *The Art of Reception*, J. Bracker and A.-K. Hubrich, Eds., Mar. 2021, pp. 371–399.
- [7] M. Thlithi, C. Barras, J. Piquier, and T. Pellegrini, "Singer diarization: application to ethnomusicological recordings," in *Proc. of the 5th International Workshop on Folk Music Analysis (FMA 2015)*, Jun. 2015, pp. 124–125.
- [8] X. Anguera, S. Bozonnet, N. Evans, C. Fredouille, G. Friedland, and O. Vinyals, "Speaker diarization: A review of recent research," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 2, pp. 356–370, Feb. 2012.
- [9] T. J. Park, N. Kanda, D. Dimitriadis, K. J. Han, S. Watanabe, and S. Narayanan, "A review of speaker diarization: Recent advances with deep learning," *Computer Speech & Language*, vol. 72, no. C, pp. 1–34, Mar. 2022.
- [10] Y. Fujita, N. Kanda, S. Horiguchi, Y. Xue, K. Nagamatsu, and S. Watanabe, "End-to-end neural speaker diarization with self-attention," in *Proc. of the 2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, Dec. 2019, pp. 296–303.
- [11] A. Mesaros and T. Virtanen, "Automatic recognition of lyrics in singing," *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2010, no. 1, pp. 1–11, Feb. 2010.
- [12] X. Gao, C. Gupta, and H. Li, "Automatic lyrics transcription of polyphonic music with lyrics-chord multi-task learning," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 2280–2294, Jul. 2022.
- [13] D. Yang and W.-S. Lee, "Music emotion identification from lyrics," in *Proc. of the 11th IEEE International Symposium on Multimedia*, Dec. 2009, pp. 624–629.
- [14] Y. E. Kim, E. M. Schmidt, R. Migneco, B. G. Morton, P. Richardson, J. Scott, J. A. Speck, and D. Turnbull, "Music emotion recognition: A state of the art review," in *Proc. of the 11th International Society for Music Information Retrieval Conference (ISMIR 2010)*, vol. 86, Aug. 2010, pp. 937–952.
- [15] C. Nithin and J. Cheriyan, "A novel approach to automatic singer identification in duet recordings with background accompaniments," in *Proc. of the 2014 Annual International Conference on Emerging Research Areas: Magnetics, Machines and Drives (AICERA/iCMMMD)*, Jul. 2014, pp. 1–6.
- [16] H. Fujihara, M. Goto, T. Kitahara, and H. G. Okuno, "A modeling of singing voice robust to accompaniment sounds and its application to singer identification and vocal-timbre-similarity-based music information retrieval," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 3, pp. 638–648, Mar. 2010.
- [17] M. Goto, T. Saitou, T. Nakano, and H. Fujihara, "Singing information processing based on singing voice modeling," in *Proc. of the 2010 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Mar. 2010, pp. 5506–5509.

- [18] F. C. Heilbron, V. Escorcia, B. Ghanem, and J. C. Niebles, "ActivityNet: A large-scale video benchmark for human activity understanding," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2015.
- [19] S. Abu-El-Haija, N. Kothari, J. Lee, P. Natsev, G. Toderici, B. Varadarajan, and S. Vijayanarasimhan, "YouTube-8M: A large-scale video classification benchmark," *arXiv [cs.CV] 1609.08675*, Sep. 2016.
- [20] D. Uthus, G. Tanzer, and M. Georg, "YouTube-ASL: A large-scale, open-domain american sign language-english parallel corpus," in *37th Conference on Neural Information Processing Systems (NeurIPS 2023) Track on Datasets and Benchmarks*, Jun. 2023.
- [21] J. F. Gemmeke, D. P. W. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. Channing Moore, M. Plakal, and M. Ritter, "Audio Set: An ontology and human-labeled dataset for audio events," in *Proc. of the 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Mar. 2017, pp. 776–780.
- [22] J. Valk and T. Alumäe, "VoxLingua107: A dataset for spoken language recognition," in *Proc. of the 2021 IEEE Spoken Language Technology Workshop (SLT)*, Jan. 2021, pp. 652–658.
- [23] S. Takamichi, L. Kürzinger, T. Saeki, S. Shiota, and S. Watanabe, "JTubeSpeech: corpus of japanese speech collected from YouTube for speech recognition and speaker verification," *arXiv [cs.SD] 2112.09323*, Dec. 2021.
- [24] X. Li, S. Takamichi, T. Saeki, W. Chen, S. Shiota, and S. Watanabe, "Yodas: Youtube-Oriented dataset for audio and speech," in *Proc. of the 2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, Dec. 2023, pp. 1–8.
- [25] A. Watanabe, S. Takamichi, Y. Saito, W. Nakata, D. Xin, and H. Saruwatari, "Coco-Nut: Corpus of Japanese utterance and voice characteristics description for prompt-based control," in *Proc. of the 2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, Sep. 2023.
- [26] A. Noulas, G. Englebienne, and B. J. A. Krose, "Multimodal speaker diarization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 1, pp. 79–93, Mar. 2011.
- [27] National Institute of Standards and Technology (NIST), "The 2009 (RT-09) rich transcription meeting recognition evaluation plan," Tech. Rep., 2009.
- [28] J. Barker, S. Watanabe, E. Vincent, and J. Trmal, "The fifth 'CHiME' speech separation and recognition challenge: Dataset, task and baselines," in *Proc. of the Interspeech 2018*, Mar. 2018, pp. 1561–1565.
- [29] A. Plaquet and H. Bredin, "Powerset multi-class cross entropy loss for neural speaker diarization," in *Proc. of the Interspeech 2023*, 2023, pp. 3222–3226.
- [30] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *Proc. of the 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 2018, pp. 5329–5333.
- [31] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao, J. Wu, L. Zhou, S. Ren, Y. Qian, Y. Qian, J. Wu, M. Zeng, X. Yu, and F. Wei, "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, Oct. 2022.
- [32] B. Desplanques, J. Thienpondt, and K. Demuynck, "ECAPA-TDNN: Emphasized channel attention, propagation and aggregation in TDNN based speaker verification," in *Proc. of the Interspeech 2020*, ISCA, Oct. 2020.
- [33] N. Dawalatabad, M. Ravanelli, F. Grondin, J. Thienpondt, B. Desplanques, and H. Na, "ECAPA-TDNN embeddings for speaker diarization," *arXiv [eess.AS] 2104.01466*, Apr. 2021.
- [34] T. L. Nguyen, B. T. Ta, D. Van Hai, T. A. X. Tran, and N. M. Le, "Speaker diarization for Vietnamese conversations using deep neural network embeddings," in *Proc. of the 2022 IEEE Ninth International Conference on Communications and Electronics (ICCE)*, Jul. 2022, pp. 300–305.
- [35] M. Ravanelli, T. Parcollet, P. Plantinga, A. Rouhe, S. Cornell, L. Lugosch, C. Subakan, N. Dawalatabad, A. Heba, J. Zhong, J.-C. Chou, S.-L. Yeh, S.-W. Fu, C.-F. Liao, E. Rastorgueva, F. Grondin, W. Aris, H. Na, Y. Gao, R. D. Mori, and Y. Bengio, "SpeechBrain: A general-purpose speech toolkit," *arXiv [eess.AS] 2106.04624*, Jun. 2021.
- [36] A. Défossez, "Hybrid spectrogram and waveform source separation," in *Proc. of the ISMIR 2021 Music Demixing Workshop*, Nov. 2021, pp. 1–11.
- [37] S. Rouard, F. Massa, and A. Défossez, "Hybrid transformers for music source separation," in *Proc. of the 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Jun. 2023, pp. 1–5.
- [38] A. Hotho, A. Nürnberger, and G. Paaß, "A brief survey of text mining," *Journal for Language Technology and Computational Linguistics*, vol. 20, no. 1, pp. 19–62, Jul. 2005.
- [39] J. Shi and J. Malik, "Normalized cuts and image segmentation," *IEEE transactions on pattern analysis and machine intelligence*, vol. 22, no. 8, pp. 888–905, 2000.



- [40] S. Horiguchi, Y. Fujita, S. Watanabe, Y. Xue, and K. Nagamatsu, “End-to-end speaker diarization for an unknown number of speakers with encoder-decoder based attractors,” in *Proc. of the Interspeech 2020*, May 2020, pp. 269–273.
- [41] H. Bredin, “Pyannote.Audio 2.1 speaker diarization pipeline: Principle, benchmark, and recipe,” in *Proc. of the Interspeech 2023*, Aug. 2023.