



MusicTM-Dataset for Joint Representation Learning Among Sheet Music, Lyrics, and Musical Audio

Donghuo Zeng^(✉), Yi Yu, and Keizo Oyama

National Institute of Informatics, SOKENDAI, Tokyo, Japan
{zengdonghuo,yiyi,oyama}@nii.ac.jp

Abstract. This work present a music dataset named MusicTM-Dataset, which is utilized in improving the representation learning ability of different types of cross-modal retrieval (CMR). Little large music dataset including three modalities is available for learning representations for CMR. To collect a music dataset, we expand the original musical notation to synthesize audio and generated sheet-music image, and build musical notation based sheet-music image, audio clip and syllable-denotation text as fine-grained alignment, such that the MusicTM-Dataset can be exploited to receive shared representation for multi-modal data points. The MusicTM-Dataset presents 3 kinds of modalities, which consists of the image of sheet-music, the text of lyrics and synthesized audio, their representations are extracted by some advanced models. In this paper, we introduce the background of music dataset and express the process of our data collection. Based on our dataset, we achieve some basic methods for CMR tasks. The MusicTM-Dataset are accessible in <https://github.com/dddzeng/MusicTM-Dataset>.

Keywords: MusicTM-Dataset · MIR · Canonical correlation analysis

1 Introduction

Music data is getting readily accessible in digital form online, which brings difficult to manage the music from a large amount of personal collection. It highly relies on the music information retrieval to retrieve the right data information for users. In recent years, machine learning or deep learning based methods has become increasing prevailing in music information retrieval [1–8] and has played an essential role in MIR.

This paper concentrates on content music MIR by learning semantic concepts across different music modalities for MIR, as shown in the Fig. 1. For instance, when we play music audio, we want to find what is the corresponding sheet music and which lyrics is correct, by learning two kinds of relationship in audio-sheet music and audio-lyrics. Such kinds of relationship obtained from content-based representation by learning the alignment across two modalities in the shared

latent subspace without introducing any users' information. The unsupervised representation learning method ensures the system can allow users to find the right music data modalities with the other data modalities as query.

The major challenge of unsupervised representation learning for different music modalities is the modality gap. Representation learning for two music data modalities such as audio-lyrics [9–11], audio-sheet music [12, 13], have become increasingly in the CMR task to bridge the modality gap. In previous works, classic CCA and CCA-variant methods [14, 15] are popular in representation learning between two music data modalities, through finding linear or nonlinear transformation to optimize the correlation between two data modalities in the shared latent subspace. With the success of Deep Neural Network (DNN) in representation learning, DNN is also helpful for learning joint representation for cross-modal tasks [16], for example, attention network [12] applies a soft-attention mechanism for the audio branch to learn the relationship between sheet music and audio, which solves the problem that the music recordings easily brings about the global and local time deviations.

However, representation learning for two modalities is still not enough to achieve the music information retrieval, when we apply one data modality as query to retrieve other two different data modalities. The existing dataset normally applied in learning correlation between two modalities in a shared space. The paper [13] collect a dataset contains an alignment between sheet music and music audio, which explores music audio to find the corresponding sheet music snippets. [17] apply a lyrics and audio paired dataset to align lyrics to audio. In this paper, we collect a new music dataset including three music data modalities. In particular, sheet music and audio are generated from music notes by music generation tools, the syllable-level lyrics and music notes are fine-grained alignment. Three major contributions of this paper have achieved in the following aspects: 1) we collect a fine-grained alignment across three music data modalities, which is useful for representation learning methods to obtain high-level feature for music CMR tasks. 2) we release experimental results of some baselines such as CCA and Generalized CCA on our MusicTM-Dataset. 3) The performance of Generalized CCA surpasses the CCA on audio-sheet music CMR task, which shows that the mapping all the three data modalities into a shared latent subspace can be better than mapping them into two shared latent subspace for audio-sheet music cross-modal retrieval.

The rest parts are arranged as follows. Some existing related works show in Sect. 2. In Sect. 3, we explain the detail of our data collection, feature representations and the metrics we applied on our experiment in Sect. 4. Section 5 makes a conclusion of the whole paper.




Fig. 1. The framework of representation learning for music information retrieval.

2 Related Works

2.1 Audio and Lyrics

Recently, the study of automatic audio-lyrics alignment techniques is getting trendy. The aim of the topic is to estimate the relation between audio and lyrics, such as temporal relation [18], deep sequential correlation [19]. [17] establishes audio-lyrics alignment based on a hidden Markov model speech recognizer, in particular, the lyrics input is to create a language model and apply the Viterbi method to link the audio and the lyrics. Synchronizing lyrics information with an audio recording is an important music application. [20] presents an approach for audio-lyric alignment by matching the vocal track and the synthesized speech.

2.2 Sheet Music and Audio

The popular problem of correlation learning between sheet music and audio is to establish the relevant linking structures between them. In [21], it aims to establish linking the regions of sheet music to the corresponding piece in an audio of the same clip. [22] bring forwards an multi-modal convolutional neural network, by taking an audio snippet as input to find the relevant pixel area in sheet music image. However, the global and local tempo deviations in music recordings will influence the performance of the retrieval system in the temporal context. To address that, [23] introduces an additional soft-attention mechanism on audio modality. Instead of correlation learning with high-level representations, [13] matches music audio to sheet music directly, the proposed method learns shared embedding space for short snippet of music audio and the corresponding piece in sheet music.

2.3 Lyrics and Sheet Music

Learning the correlation between lyrics and sheet music is a challenging research issue, which requires to learning latent relationship with high-level representations. The automatic composition techniques are considerable for upgrading music applications. [24] proposed a novel deep generative model LSTM-GAN to learn the correlation in lyrics and melody for generation task. Similarly, [25] presents an approach that is used to generate music song from a Japanese lyrics. [26] introduces a novel language model that can generate lyrics from a given sheet music. [27] presents an better query in using lyrics and melody, which take advantage of extra lyrics information by linking the scores from pitch-based lyrics and melody recognition. Accept that, “singing voice,” which is for generating singing voice has been drawing attention in the last years, [28] explores a novel model that the singing voice generation with no consideration of pre-assigned melody and lyrics.

3 Dataset and Metrics

This section presents the motivation and contribution of our data collection. Moreover, also the process of dataset collection applied in our experiments and the data feature extraction are discussed. In the end, we show all the evaluation metrics applied to leverage our models.

3.1 Dataset Collection

Figure 2 shows a few examples of MusicTM-Dataset we applied, including the spectrum of music audio with Librosa library¹, word-level lyrics, and sheet music with Lilypond technique².

The available music dataset with three modalities, which can be applied in music information retrieval based on the high-level semantic features is rarely reported. We try to learn aligned representation for sheet music images, music audio, and lyrics because they frequently appear in the music data collection. We follow the work [24] to collect our music dataset by extending two modalities (lyrics and music notes) to three modalities: sheet music, audio, and lyrics.

In [24] presents a music dataset that a music is represented by lyrics and music notes. The lyrics is parsed as syllable-level collection, such as the lyrics: ‘Listen to the rhythm of fall ...’ will parse as ‘Lis ten to the rhy thm of fall’. A music note is a ternary structure that includes three attributions: pitch, duration, and rest. The pitch is a frequency-related scale of sounds, for example, piano keys MIDI number ranges from 21 to 108, each MIDI number corresponds to a pitch number, such as MIDI number ‘76’ represents pitch number ‘E5’. Duration in music notes denotes the time of the pitch, for example, a pitch number ‘E5’ with its duration 1.0, means this music note will last 0.5 s in the playing. The rest of

¹ <https://librosa.org/doc/latest/index.html>.

² <http://lilypond.org/>.




Fig. 2. Examples of three data modalities in MusicTM-Dataset.

the pitch is the intervals of silence between two adjacent music notes, which share the same unit with duration. The dataset used for the melody generation from lyrics, to consider the time-sequence information in the pairs, the syllable-level lyrics and music notes are aligned by pairing a syllable and a note.

The initial pre-processing for our dataset is to get the beginning of music notes and corresponding syllables. In our MusicTM-Dataset collection, we adopted the same method to get the first 20 notes as a sample and ensure the syllable-level lyrics corresponding can be kept. Moreover, we removed the samples if existing the rest attributes of the note are longer than 8 (about four seconds).

Music audio and sheet music are separately created from music notes that matches our purpose of musical multimodal building. We use syllable-level lyrics and notes to create the pairs of sheet and audio by some high-quality technologies. All the music data modalities contain temporal structure information, which motivates us to establish fine-grained alignment across different modalities, as seen in Fig. 3. In detail, the syllable of lyrics, the audio snippet, and sheet music fragment generated from music notes are aligned.

Music audio is also music sound transmitted in signal form. We add piano instrument in the music channel to create new midi files, and synthesize audios with TiMidity++ tool³.

Sheet music is created by music note with Lilypond tools. Lilypond is a compiled system that runs on a text file describing the music. The text file may contain music notes and lyrics. The output of Lilypond is sheet music which can be viewed as an image. Lilypond is like a programming language system, music notes are encoded with letters and numbers, and commands are entered with backslashes. It can combine melody with lyrics by adding the “\addlyrics” command. In our MusicTM-Dataset, sheet music (visual format) for one note and entire sheet music (visual format) for 20 notes are created respectively. Accordingly, each song has single note-level and sequential note-level (sheet fragment) visual formats.

³ <http://timidity.sourceforge.net/>.




Fig. 3. An example of fine-grained alignment across three modalities.




Fig. 4. The audio feature extraction process with vggish model

3.2 Feature Extraction

This section will explain the feature extraction for music multimodal data.

Audio Feature Extraction. Generally, audio signal is used for audio feature extraction, which plays the main role in speech processing [29,30], music genre classification [31], and so on. Here, we present a typical model for audio feature extraction, the supervised trained model Vggish. The detailed process of feature extraction can be seen in Fig. 4. Firstly, we resample audio waveform 16 kHz mono, then calculate a spectrogram. Secondly, in order to obtain a stable log mel spectrogram, it is computed by exploring log. Finally, resampling the feature into (125,64) format, then applying pre-trained model to extract feature and use PCA model to map it into 128-dimensional.

Sheet Music Feature Extraction. Different from other image feature extraction, our feature extraction of sheet music image tries to catch pitches and the segments. In this paper, our information extraction of sheet music has two levels, pitch detection, and semantic segments. We apply the ASMCMR [32] model trained in audio-sheet retrieval tasks, which learns the correlation between audio clips and corresponding sheet snippets. In our work, the shape of extracted note-level feature and sheet snippet-level features are (100, 32) and (32,) respectively.

Table 1. Statistics of MusicTM-Dataset applied in our experiments

Modality	Feature extractor	Dimension	Number
Audio	Vggish	(20, 128)	14,454
Lyrics	Skip-gram	(20, 20)	14,454
Sheet music	Lilypond&ASMCMR	(20, 100, 32)	14,454

Lyrics Feature Extraction. We follow [24] to keep the alignment between syllable and note by representing lyrics in the form of syllable and word level. The syllable-level feature extracted with the syllable skip-gram model, the word-level feature extracted with the word skip-gram model used in [24]. These two pre-trained skip-gram models are trained on all the lyrics data, which applied in a regression task with SGD optimization. The input of syllable-level skip-gram model is a sequence of syllables in a sentence, while the input of word-level model is a word unit sequence in the sentence. The output of the syllable-level and word-level skip-gram model is 20-dimensional embedding for each syllable and word, respectively.

The overall statistics of our music data are shown in Table 1. We divided the dataset into 3 parts as training, validation, and testing set by 70%, 15%, and 15%. The number of training, validation, and testing set are 13,535, 2800, and 2800 respectively.

3.3 Evaluation Metric

To evaluate some baselines on our dataset, we apply some standard evaluation from the work [33] for unsupervised learning based cross-modal retrieval. R@K (Recall at K, here we set K as 1, 5, and 10) is to compute correct rate that is the percentage of retrieved items corresponding to the query in the top-K of rank list. For instance, R@1 calculate the percentage of sample appear in the first item of retrieved list. In order to further evaluate our collected dataset with some baselines, we also apply the Median Rank and Mean Rank to compute the mean and median rank of all the correct results.

4 Experiments

Table 2. The performance of multimodal information retrieval on MusicTM-Dataset.

audio2lyrics retrieval					
Methods	R@1	R@5	R@10	MedR	MeanR
Random rank [34]	0.028	0.055	0.076	7312.0	7257.2
CCA [35]	0.306	0.350	0.353	423.0	639.4
GCCA [36]	0.040	0.074	0.093	770.0	881.1
lyrics2audio retrieval					
Random rank	0.027	0.055	0.076	7316.0	7257.3
CCA	0.304	0.349	0.354	427.0	639.3
GCCA	0.039	0.078	0.095	774.0	881.6
sheet music2lyrics retrieval					
Random rank	0.027	0.055	0.075	7311.0	7257.3
CCA	0.093	0.172	0.203	524.0	708.7
GCCA	0.089	0.0142	0.167	573.0	770.5
lyrics2sheet music retrieval					
Random rank	0.027	0.055	0.077	7313.0	7257.4
CCA	0.093	0.168	0.198	522.0	709.0
GCCA	0.098	0.014	0.168	578.0	769.8
audio2sheet music retrieval					
Random rank	0.028	5.57	7.50	7310.0	7257.2
CCA	0.303	0.349	0.353	341.0	596.5
GCCA	0.358	0.403	0.414	271.0	382.8
sheet music2audio retrieval					
Random rank	0.026	0.055	0.075	7310.0	7257.4
CCA	0.300	0.350	0.354	332.0	596.1
GCCA	0.362	0.407	0.415	271.0	381.3

4.1 Baselines

CCA can be seen as the method that aims at finding linear transforms for two sets of variables in order to optimize the relation between the projections of the variable sets into a shared latent subspace. Consider two variables from two data modalities $X \in R^{D_x}$ and $Y \in R^{D_y}$ with zero mean and the two paired data sets $S_x = \{x_1, x_2, \dots, x_n\}$ and $S_y = \{y_1, y_2, \dots, y_n\}$. $W_x \in R^{D_x}$ and $W_y \in R^{D_y}$ as the directions that linearly map the two set into a shared latent subspace, such that the relation between the projection of S_x and S_y on W_x and W_y is optimized.

$$\rho = \arg \max_{(W_x, W_y)} \frac{W_x^T \Sigma_{xy} W_y}{\sqrt{W_x^T \Sigma_{xx} W_x \cdot W_y^T \Sigma_{yy} W_y}} \quad (1)$$

where ρ is the correlation, Σ_{xx} and Σ_{yy} denote the variance-covariance matrix of S_x , S_y , respectively and Σ_{xy} represents the cross-covariance matrix.

Generalized CCA [36] can be viewed as an extension method of CCA, which aims to solve the limitation on the number of data modalities. The objective function in Eq. 2, which focuses on finding a shared representation G for K different data modalities.

$$\text{minimize}_{(W_k, G)} = \sum_{k=1}^K \|G - W_k^T X_k\|_F^2 \quad (2)$$

where K is the size of data points, and X_k is a matrix for k^{th} data modality. Similar to CCA, GCCA is to find linear transformation for different data modalities to optimize the correlation within them.

4.2 Results

In Table 2, when learning the correlation between two data modalities with CCA method, the correlation of audio-lyrics and audio-sheet music can get more than 30% of R@1, which illustrates the dataset can be learned for cross-modal retrieval task. Specifically, in comparison with CCA and RANDOM, GCCA will have a big improvement in the performance of audio-sheet music cross-modal retrieval. In detail, compared with CCA method, 5.46%, 5.39%, 6.06%, 70, and 213.68 improved in R@1, R@5, R@10, MedR, and MeanR for music audio as the query to retrieve the correct sheet music; 6.16%, 5.65%, 6.1%, 61, and 214.8 improved in R@1, R@5, R@10, MedR, and MeanR for sheet music as the query to retrieve the correct music audio. However, GCCA will decrease the performance of audio-lyrics cross-modal retrieval and achieve a similar performance of sheet music-lyrics cross-modal retrieval.

The results show that the learned representation with GCCA for sheet image, lyrics, and music audio can raise the relation of sheet music and music audio. However, such representations drop the correlation between music audio and lyrics and their correlation between sheet music image and lyrics will almost stay the same as CCA method, which learns the representation in the shared subspace without involving lyrics data. The results prove our hypothesis can be accepted that the sheet music and music audio are created by music notes, so the correlation between audio and sheet music will be close. The lyrics and music note from original dataset exist alignment between each other, the correlation between the two can be learned. In this case, the correlation between audio and lyrics reflects the correlation between audio and music note, however, the correlation between sheet music and lyrics seems hard to learn.

In visualization of the the position of sheet music, lyrics, and music audio in CCA and GCCA subspace, as shown in Fig. 5. GCCA seems to pull audio and sheet music while pushing the audio and lyrics compared with the CCA subspace. This motivates us to propose a new advanced model that can improve three couples of cross-modal retrieval tasks in a shared latent subspace as the GCCA subspace achievement in the future.




Fig. 5. The general paradigm of MusicTM-Dataset with two different models (CCA, GCCA)

5 Conclusion

This paper presents a MusicTM-Dataset that consists of three different data modalities and there is fine-grained alignment across the modalities. The dataset can be easily extended to different researches, we report the performance of some baselines on our MusicTM-Dataset, which allows the results of the following research to be compared. Instead of applying CCA to learn shared latent subspace for every two modalities, GCCA learns the correlation of three modalities in one shared latent subspace. The performance of audio-sheet music can be improved and the performance of audio-lyrics cross-modal retrieval is quite similar but the performance of lyrics-sheet music cross-modal retrieval will be decreased. In theory, we want to develop a new architecture that will improve the performance of multimodal information retrieval across different modalities.

Acknowledgements. The JSPS Grant for SR financed this work, which is under Grant No. 19K11987.

References

1. Eyben, F., Böck, S., Schuller, B., Graves, A.: Universal onset detection with bidirectional long-short term memory neural networks. In: Proceedings of the 11th International Society for Music Information Retrieval Conference, ISMIR, Utrecht, The Netherlands, pp. 589–594 (2010)
2. Hamel, P., Eck, D.: Learning features from music audio with deep belief networks. In: ISMIR, Utrecht, The Netherlands, vol. 10, pp. 339–344 (2010)
3. Zhou, X., Lerch, A.: Chord detection using deep learning. In: Proceedings of the 16th ISMIR Conference, vol. 53, p. 152 (2015)
4. Böck, S., Krebs, F., Widmer, G.: Accurate tempo estimation based on recurrent neural networks and resonating comb filters. In: ISMIR, pp. 625–631 (2015)
5. Grill, T., Schluter, J.: Music boundary detection using neural networks on spectrograms and self-similarity lag matrices. In: 2015 23rd European Signal Processing Conference (EUSIPCO), pp. 1296–1300. IEEE (2015)

6. Choi, K., Fazekas, G., Cho, K., Sandler, M.: A tutorial on deep learning for music information retrieval. arXiv preprint [arXiv:1709.04396](https://arxiv.org/abs/1709.04396) (2017)
7. Siedenburg, K., Fujinaga, I., McAdams, S.: A comparison of approaches to timbre descriptors in music information retrieval and music psychology. *J. New Music Res.* **45**(1), 27–41 (2016)
8. Sigtia, S., Boulanger-Lewandowski, N., Dixon, S.: Audio chord recognition with a hybrid recurrent neural network. In: *ISMIR*, pp. 127–133 (2015)
9. Yu, Y., Tang, S., Raposo, F., Chen, L.: Deep cross-modal correlation learning for audio and lyrics in music retrieval. *ACM Trans. Multimedia Comput. Commun. Appl. (TOMM)* **15**(1), 1–16 (2019)
10. Kruspe, A.M., Fraunhofer, I.D.M.T.: Retrieval of textual song lyrics from sung inputs. In: *INTERSPEECH*, pp. 2140–2144 (2016)
11. Kruspe, A.M., Goto, M.: Retrieval of song lyrics from sung queries. In: *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 111–115. IEEE (2018)
12. Balke, S., Dorfer, M., Carvalho, L., Arzt, A., Widmer, G.: Learning soft-attention models for tempo-invariant audio-sheet music retrieval. arXiv preprint [arXiv:1906.10996](https://arxiv.org/abs/1906.10996) (2019)
13. Matthias, D., Hajič Jr., J., Arzt, A., Frostel, H., Widmer, G.: *Transactions of the International Society for Music Information Retrieval* **1**(1) (2018)
14. Dorfer, M., Arzt, A., Widmer, G.: Towards end-to-end audio-sheet-music retrieval. arXiv preprint [arXiv:1612.05070](https://arxiv.org/abs/1612.05070) (2016)
15. Dorfer, M., Schlüter, J., Vall, A., Korzeniowski, F., Widmer, G.: End-to-end cross-modality retrieval with CCA projections and pairwise ranking loss. *Int. J. Multimedia Inf. Retrieval* **7**(2), 117–128 (2018)
16. Yu, Y., Tang, S., Aizawa, K., Aizawa, A.: Category-based deep CCA for fine-grained venue discovery from multimodal data. *IEEE Trans. Neural Netw. Learn. Syst.* **30**(4), 1250–1258 (2018)
17. Mauch, M., Fujihara, H., Goto, M.: Integrating additional chord information into hmm-based lyrics-to-audio alignment. *IEEE Trans. Audio Speech Lang. Process.* **20**(1), 200–210 (2011)
18. Fujihara, H., Goto, M.: Lyrics-to-audio alignment and its application. In: Müller, M., Goto, M., Schedl, M. (eds.) *Multimodal Music Processing. Dagstuhl Follow-Ups*, vol. 3, pp. 23–36. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany (2012)
19. Yu, Y., Tang, S., Raposo, F., Chen, L.: Deep cross-modal correlation learning for audio and lyrics in music retrieval. *CoRR*, abs/1711.08976 (2017)
20. Lee, S.W., Scott, J.: Word level lyrics-audio synchronization using separated vocals. In: *2017 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2017, New Orleans, LA, USA, 5–9 March 2017*, pp. 646–650. IEEE (2017)
21. Thomas, V., Fremerey, C., Müller, M., Clausen, M.: Linking sheet music and audio - challenges and new approaches. In: Müller, M., Goto, M., Schedl, M. (eds.) *Multimodal Music Processing. Dagstuhl Follow-Ups*, vol. 3, pp. 1–22. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany (2012)
22. Dorfer, M., Arzt, A., Widmer, G.: Towards score following in sheet music images. In: Mandel, M.I., Devaney, J., Turnbull, D., Tzanetakis, G. (eds.) *Proceedings of the 17th International Society for Music Information Retrieval Conference, ISMIR 2016, New York City, United States, 7–11 August 2016*, pp. 789–795 (2016)

23. Balke, S., Dorfer, M., Carvalho, L., Arzt, A., Widmer, G.: Learning soft-attention models for tempo-invariant audio-sheet music retrieval. In: Flexer, A., Peeters, G., Urbano, J., Volk, A. (eds.) Proceedings of the 20th International Society for Music Information Retrieval Conference, ISMIR 2019, Delft, The Netherlands, 4–8 November 2019, pp. 216–222 (2019)
24. Yu, Y., Canales, S.: Conditional LSTM-GAN for melody generation from lyrics. arXiv preprint [arXiv:1908.05551](https://arxiv.org/abs/1908.05551) (2019)
25. Fukayama, S., Nakatsuma, K., Sako, S., Nishimoto, T., Sagayama, S.: Automatic song composition from the lyrics exploiting prosody of the Japanese language. In: Proceedings of the 7th Sound and Music Computing Conference (SMC), pp. 299–302 (2010)
26. Watanabe, K., Matsubayashi, Y., Fukayama, S., Goto, M., Inui, K., Nakano, T.: A melody-conditioned lyrics language model. In: Walker, M.A., Ji, H., Stent, A. (eds.) Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, 1–6 June 2018, vol. 1 (Long Papers), pp. 163–172. Association for Computational Linguistics (2018)
27. Wang, C.-C., Roger Jang, J.-S., Wang, W.: An improved query by singing/humming system using melody and lyrics information. In: Stephen Downie, J., Veltkamp, R.C. (eds.) Proceedings of the 11th International Society for Music Information Retrieval Conference, ISMIR 2010, Utrecht, Netherlands, 9–13 August 2010, pp. 45–50. International Society for Music Information Retrieval (2010)
28. Liu, J.-Y., Chen, Y.-H., Yeh, Y.-C., Yang, Y.-H.: Score and lyrics-free singing voice generation. arXiv preprint [arXiv:1912.11747](https://arxiv.org/abs/1912.11747) (2019)
29. Watanabe, S., Hori, T., Karita, S., Hayashi, T., Nishitoba, J., Unno, Y., Soplin, N.E.Y., Heymann, J., Wiesner, M., Chen, N., et al.: ESPnet: end-to-end speech processing toolkit. arXiv preprint [arXiv:1804.00015](https://arxiv.org/abs/1804.00015) (2018)
30. Lambert, C., Kormos, J., Minn, D.: Task repetition and second language speech processing. *Stud. Second Lang. Acquisition* **39**(1), 167–196 (2017)
31. Kobayashi, T., Kubota, A., Suzuki, Y.: Audio feature extraction based on sub-band signal correlations for music genre classification. In: 2018 IEEE International Symposium on Multimedia (ISM), pp. 180–181. IEEE (2018)
32. Dorfer, M., Arzt, A., Widmer, G.: Learning audio-sheet music correspondences for score identification and offline alignment. In: Proceedings of the 18th International Society for Music Information Retrieval Conference, ISMIR 2017, Suzhou, China, 23–27 October 2017, pp. 115–122 (2017)
33. Otani, M., Nakashima, Y., Rahtu, E., Heikkilä, J., Yokoya, N.: Learning joint representations of videos and sentences with web image search. In: Computer Vision - ECCV 2016 Workshops - Amsterdam, The Netherlands, 8–10 and 15–16 October 2016, Proceedings, Part I, pp. 651–667 (2016)
34. Zeng, D., Yu, Y., Oyama, K.: Unsupervised generative adversarial alignment representation for sheet music, audio and lyrics. arXiv preprint [arXiv:2007.14856](https://arxiv.org/abs/2007.14856) (2020)
35. Hardoon, D.R., Szedmák, S., Shawe-Taylor, J.: Canonical correlation analysis: an overview with application to learning methods. *Neural Comput.* **16**(12), 2639–2664 (2004)
36. Tenenhaus, A., Tenenhaus, M.: Regularized generalized canonical correlation analysis. *Psychometrika* **76**(2), 257 (2011)