

LISTEN, READ, AND IDENTIFY: MULTIMODAL SINGING LANGUAGE IDENTIFICATION OF MUSIC

Keunwoo Choi and Yuxuan Wang
ByteDance

{keunwoo.choi, wangyuxuan.11}@bytedance.com

ABSTRACT

We propose a multimodal singing language classification model that uses both audio content and textual metadata. LRID-Net, the proposed model, takes an audio signal and a language probability vector estimated from the metadata and outputs the probabilities of the target languages. Optionally, LRID-Net is facilitated with modality dropouts to handle a missing modality. In the experiment, we trained several LRID-Nets with varying modality dropout configuration and tested them with various combinations of input modalities. The experiment results demonstrate that using multimodal input improves performance. The results also suggest that adopting modality dropout does not degrade the performance of the model when there are full modality inputs while enabling the model to handle missing modality cases to some extent.

1. INTRODUCTION

Singing language identification of music (SLID) is a classification task of labeling tracks by the languages used in lyrics. Language information is essential for music discovery and recommendation systems as lyrics play a crucial role in the music listening experience [1]. In reality, despite its importance, language information is not always available or accurate, even for established, large-scale music streaming services. Such a situation has motivated researchers to develop language identification models using the most fundamental music data – the audio content itself.

There have been various audio-based approaches in SLID. Most of them utilize traditional machine learning classifiers and audio features, similar to early methods for spoken language identification [2, 3]. The models in [4] and [5] consist of vocal/non-vocal segmentation, feature vector quantization, and a simple codebook-based language model. More recently, *i*-vector [6], a popular feature vector for speech-related tasks, was used in [7, 8], combined with a support vector machine classifier. A vocal source separation technique that is based on a spatial property of stereo music signals was added in a model [9], although it did not improve its performance in the experi-

ment. More recently, the model in [10] uses a modern vocal separation technique [11] and a one-dimensional deep convolutional neural network. In these approaches, a vocal source separation module is applied to input signals to extract relevant features with a cleaner signal, i.e., with a less amount of accompaniments.

There also have been models that are based on non-audio modalities. The model in [12] is designed to classify music videos into language categories by taking visual features such as histograms of oriented gradients along with basic audio features such as MFCCs. In the experiment, adding video features improved the accuracy of the model from 44.7% to 47.8% in a 25-language classification task. Another model in [13] uses language estimation of track title and album name and showed a comparable performance to their in-house audio-based classifier. Notably, an internal music representation called track vector, which is estimated using music listening history data, showed the highest feature relevance – 0.97 – in their experiment. It re-emphasizes that there is a strong connection between music listening preference and singing language of music.

It is noteworthy that unfortunately, none of the mentioned works is reproducible [4, 5, 7–10, 12, 13] and there has not been any benchmark in SLID. All the previous works rely on private datasets and only little details such as target languages and the number of tracks are known. This is because of a lack of datasets – there has not been a publicly available music language classification dataset until very recently [14]. Some music tagging datasets may be considered as alternatives if they include language tags; for example, the million song dataset has language labels [15]. But their tag popularity is merely at 116th (‘german’), 135th (‘english’), or below, which are often excluded in a prevalent problem formulation such as top-50 classification [16] to suppress noise during training and evaluation [17].

In this paper, we introduce LRID-Net – **Listen, Read, and Identify-Network**, a model that takes audio and textual metadata (track title, album name, and artist name) to identify singing language. LRID-Net is based on a combination of a deep convolutional neural network, MLPs, and modality dropouts as explained in Section 4. We provide brief information about Music4All [14], the dataset that we use, in Section 3. In Section 5, we present our experiment results including the performance of LRID-Net, a comparison of modalities, and verification of modality dropout.



There are three main contributions of our work.

Contribution 1: LRID-Net is the first work in SLID that takes advantage of audio and text inputs, presumably the two most accessible forms of music data.

Contribution 2: We present the first reproducible work in SLID by using a public dataset. This enables a rigorous benchmark to be followed, which is necessary for progress in modern machine learning research.

Contribution 3: LRID-Net has flexibility in its input data form – it is designed to work with missing modalities by adopting *modality dropouts*. This flexibility makes LRID-Net a highly pragmatic solution in a real-world scenario.

2. PROBLEM FORMULATION

We define our problem as a *singing language identification using audio content and metadata*. It is reasonable to expect that in many practical use-cases, (some of) three selected types of metadata – track title, album name, and artist name – would be accessible. For example, music tracks shared on online streaming services usually include all of them. We exclude other types of data such as visual features and pre-computed track vectors despite their benefits shown in [12] and [13], respectively, because their availability is limited for a subset of commercial music tracks even for those in industry who have an access to a large-scale proprietary catalog.

To make our model even more practical, we consider a missing data scenario. Some or all of the metadata can be easily missing. For example, indie music tracks shared online (e.g., SoundCloud or Jamendo) usually do not include any album name, or they can exist but should be considered missing since titles can be blank or consists of numbers and/or special characters only, not providing any linguistic information. The audio could be also missing or not helpful at all, for example, a segment that is input to the model may not contain any vocal part.

From a machine learning point of view, our SLID problem is a single-label multi-class classification with a multi-modal, potentially partially missing input. In fact, some lyrics are multilingual. But we assume that those cases are negligible, following the dataset we use (see Section 3 for more details).¹

3. DATASET

We use Music4All dataset [14] which includes 30-second audio clips (44,100 kHz and stereo), lyrics, and 16 other metadata such as title, album name, artist name, and Spotify identifier of 109,269 tracks. The dataset also includes language labels covering 46 languages. They are estimated from the lyrics using Langdetect,² a Python implementation of Language-Detection [18].³

¹ English words such as “Yeah”, “Hey”, and “Baby” are often used as musical expressions or interjections in non-English songs and we do not assume their existence makes a lyric multilingual.

² <https://pypi.org/project/langdetect/>

³ We noticed there are errors in the ground truth. But Langdetect is known to perform at 99% accuracy with documents, which is significant higher than the performance of the proposed audio-based models.




Figure 1: Block diagram of LRID-Net. Gray rectangular boxes indicate non-trainable modules and dotted circles are modality dropouts (Section 4.4), which are optionally applied in some of the experiments (Section 5.4).

The distribution of language labels in Music4ALL is heavily skewed from 84,103 items (English) to 1 item (Hindi, Slovak, Bulgarian, and Hebrew). We consider 11 labels – top-10 popular languages and an “others” category that includes all the other languages. In detail, there are 84,103 English tracks followed by Portuguese (7,020), Spanish (3,225), Korean (1,145), Others (1,059), French (994), Japanese (615), German (577), Polish (446), Italian (437), and Slovakian (231). There are also 9,417 instrumental tracks, but we exclude them because otherwise, the model would need to learn to perform language identification as well as instrumental track classification, which would distract our analysis.

There does not exist an official training-validation split of tracks of M4A dataset. We use an 80:20 stratified split with allocating every artist in only one of the split sets.⁴ This prevents artist-dependent information from being a confounding factor and leaking the information between training and validation sessions.

4. LRID-NET

4.1 Input Audio Preprocessing

The 44.1 kHz sampled 30-second stereo audio signals are downmixed and resampled to 22,050 Hz and converted into 128-bin log-magnitude mel-spectrograms (128 × 2580) with 1024-point FFT and 256 hop size using Kapre [19]. We call this spectrogram x_{audio} .

4.2 Input Text Preprocessing

The metadata strings are joined in an order of artist name, album name, and track title and then input to Langdetect to estimate a language probability vector. Originally, `langdetect.detect_langs()` outputs a probability distribution of 55 supported languages. However, there are some cases where the function fails to estimate probability, e.g., if the text is blank or a numeric value only. We add another dimension to indicate those exceptions. This 56-dimensional vector is called x_{lang} .

4.3 Overall model structure

LRID-Net consists of two input branches that are concatenated at a late stage of the network.

⁴ https://github.com/keunwoochoi/music4all_contrib

The audio branch, f , is a ResNet-50 with 64 base channels and outputs a 2048-channel feature map that is a size of (4×81) [20]. We choose ResNet-50 for its simplicity yet strong performance as shown in [21] for music tagging. Then, a global average pooling is applied to output a one-dimensional vector length of 2048. This procedure is represented as $s_{\text{audio}} = f(x_{\text{audio}})$.

The text branch, g , is a 3-layer MLP where each layer consists of a 128-unit fully-connected layer, a batch normalization layer, and a ReLU activation [22], i.e., $s_{\text{text}} = g(x_{\text{lang}})$.

In the output branch, h , the two outputs of the input branches are concatenated, i.e., $s_{\text{cat}} = [s_{\text{audio}}; s_{\text{text}}]$ and input to an MLP. This MLP consists of a 256-unit fully-connected hidden layer, a batch normalization layer, a ReLU activation, and an 11-unit fully-connected layer with a Softmax activation to output the language probability, i.e., $\hat{y} = h(s_{\text{cat}})$.

4.4 Modality Dropout

We use *modality dropout* which was originally introduced in [23] as *ModDrop* and is illustrated in the block diagram in Figure 1. Similar to the original dropout [23], during training time, a modality dropout module replaces its input with zeros with a probability of r , the dropout rate. There are two main differences between the original dropout and modality dropout. 1) The original dropout is applied to a part (e.g., a single node or a channel) of an input but a modality dropout module drops the whole input of a modality, i.e., an audio signal or a language probability vector. By doing so, it effectively simulates a missing modality input and lets the model learn to perform the task without the dropped input. 2) There is no $1/(1-r)$ scaling in a modality dropout when the input is not dropped. During test time, a system with LRID-Net inputs a zero vector to the model if a modality is missing.

5. EXPERIMENT AND DISCUSSION

We performed a series of experiments to demonstrate the performance and properties of LRID-Net. We use Music4All dataset [14] and process audio and text data using Kapre [19] and Langdetect [18] as detailed in Section 3 and Section 4, respectively. During training, we use Adam optimizer [24] and early stopping with a patience of 20 epochs. We do not adopt any balancing during batching and loss computation.

On the metric, we use F1-score, precision, and recall. As described in Section 3, there is a high imbalance of the number of data points of each language in the dataset. In this case, from a user perspective, a macro average can be used to represent the class-balanced performance to avoid the bias towards popular languages, too. However, *because* they are biased in the *same way* in the training and validation sets, weighted (or micro) averaging can be considered to be more suitable than macro averaging on representing how successfully the model was trained to minimize the empirical loss. Acknowledging this issue, we use both of

Text input	Precision	Recall	F1-Score
Artist Name	.323	.221	.149
Album Name	.399	.378	.284
Track Title	.450	.444	.317
Joining All	.510	.569	.429

Table 1: The macro averaged performance of Langdetect prediction with various text inputs. ‘Joining All’ represents the performance of langdetect baseline.

Text input	Precision	Recall	F1-Score
Artist Name	.600	.368	.456
Album Name	.750	.576	.652
Track Title	.766	.573	.656
Joining All	.922	.819	.857

Table 2: The weighted averaged performance of Langdetect prediction with various text inputs. ‘Joining All’ represents the performance of Langdetect baseline.

the averaging methods and focus on the performance with a suitable one depending on the context.

We present the performance of each language sorted by its occurrence count, as known as Support, to help understanding of any related trend.

5.1 Langdetect baseline

We present the result of a simple solution that is to directly use the top prediction of text-based language identification using Langdetect. This baseline approach is called *Langdetect baseline*. We also present a detailed analysis of the performance and behavior of Langdetect baseline model to deepen our understanding of the problem and the following experiment results.

Table 1 and 2 summarize the performance of Langdetect baseline based on various text inputs, showing that using all the metadata (‘Joining All’) performs the best by achieving an F1-score of 0.429 (macro average) or 0.857 (weighted average). Details of the performance with ‘Joining All’ are presented in Table 3 and its confusion matrix is illustrated in Figure 2. Note that Langdetect baseline model is not trained with our dataset, hence support (number of true items) does not affect the performance.

The F1-scores seem under-performing since Langdetect is reported to show 99% accuracy. We conjecture two reasons for this result. First, music metadata is significantly shorter than typical documents and news article, with which Langdetect was originally trained and tested, respectively. Second, there is a prevailing usage of English for artist name, album name, and track title even if the lyrics are not written in English, especially for those songs that are internationally consumed.

In detail, the precision and recall shows interesting patterns that are partly related to the second reason; There are languages where precision is significantly higher than recall (Group 1: Korean, Japanese) while precision is significantly lower than recall for some other languages (Group 2: French, German, Italian, Slovakian). We note that

Language	Precision	Recall	F1-Score	Support
English	.969	.859	.911	84,103
Portuguese	.897	.695	.783	7,020
Spanish	.676	.620	.647	3,225
Korean	.444	.003	.007	1,145
Others	.092	.700	.162	1,059
French	.367	.585	.452	994
Japanese	.895	.153	.261	615
German	.093	.773	.166	577
Polish	.695	.572	.627	446
Italian	.229	.748	.350	437
Slovakian	.252	.554	.347	231

Table 3: The performance of Langdetect prediction when all the metadata are concatenated (‘Joining all’). Support column indicates the number of items in the training set, of which distribution is preserved in the validation set.




Figure 2: The confusion matrix of Langdetect baseline model (‘Joining all’). Item counts are converted by $\log_{10}(x + 1)$.

the languages in Group 1 are only two non-European languages and suggest two explanations for their patterns. First, on high precision scores, since Korean and Japanese use completely different letter systems compared to other languages, it is very easy for Langdetect to recognize them if the metadata is written in Korean or Japanese. Second, the low recall may come from the common usage of English as mentioned earlier. In such cases, metadata-based Langdetect would never be able to correctly predict that the lyrics are written in Korean or Japanese. This is revealed in a deeper dataset analysis. Among the 1,145 Korean songs in the training set, there are 244 unique artists, out of which 241 artist names are English. Conversely, the model almost never misclassifies non-Korean or non-Japanese songs to Korean or Japanese, respectively. Unlike Group 1, we did not find any convincing explanation for the patterns of Group 2.

5.2 Single modality baselines

As additional baseline models, we show experiment results of single modality models. They are AO (audio-only)




Figure 3: Top - The performance of an text-only model (TO model). Its precision, recall, and F1-score are .526 / .415 / .422 (macro averaging) and .896 / .914 / .900 (weighted averaging). Bottom - The performance of an audio-only model (AO model). The precision, recall, and F1-score are respectively .387 / .248 / .275 (macro averaging) and .852 / .884 / .857 (weighted averaging).




Figure 4: The performance of Main model with both modality inputs. The precision, recall, and F1-score are respectively .688 / .435 / .504 (macro averaging) and .911 / .922 / .911 (weighted averaging).

model and TO (text-only) model.

First of all, we compare the TO model against Langdetect baseline. TO model achieves a comparable macro averaged F1-score (0.422, a degradation of 0.007) and a higher weighted averaged F1-score (0.900, an improvement of 0.043) as in Figure 3. Again, this result seems heavily affected by the class imbalance of the training set.

Second, as illustrated in Figure 3, a lack of a modality leads to negative effects, often critically to some languages. AO model completely fails at identifying Others, Polish, Italian, and Slovakian while TO model fails at Korean, Japanese, German, and Slovakian.

Third, as summarized in Figure 3, in every metric and averaging strategy, TO model outperformed AO model, showing the importance of using metadata. However, this does not mean audio is less useful than textual data. Acknowledging the limit of the information in the metadata discussed with Langdetect baseline in Section 5.1, the result may indicate the opposite that currently, the information from text input is almost saturated and more improvement should be based on a better audio understanding.

5.3 LRID-Net: Main Model

In this section, we introduce the experiment result of our multimodal SLID model, LRID-Net. This LRID-Net net was trained without any modality dropouts and we call this model ‘Main model’.




Figure 5: The confusion matrix of Main model. Item counts are converted by $\log_{10}(x + 1)$.

Main model shows an improvement over all the aforementioned baseline models. Compared to Langdetect baseline, it achieves a higher F1-score (+0.048) with a higher precision (+0.134) and a lower recall (-0.118) on weighted average. It also achieves a higher F1-score (+0.075) and a higher precision (+0.158), but a lower recall (-0.134) on macro average. Main model also outperforms AO model and TO model, the single modality models, in every metric and averaging strategy, emphasizing the benefit of using multimodal information.

Among languages, Main model shows low recall rates for Korean, French, Japanese, German, and Italian as shown in Figure 4. This pattern is similar to that in the results of Langdetect baseline as discussed in Section 5.1. We conjecture that a similar type of confusion may have happened in Main model, especially if the 56-dimensional text-based language probability input computed with Langdetect plays an important role (which seems true given that TO model outperforms AO model in Section 5.2). The class imbalance of the training set seems to penalize recall rates of those languages because a classification of unconfident items is likely to be biased towards the mode of the distribution of training items, i.e., English. This is shown in Figure 5.

The overall improvement of F1-score did not benefit all the languages equally. The performance of Main model is rather polarized than Langdetect baseline model. The F1-scores of 3/4 popular languages (English, Spanish, and Korean) and Others category are improved in Main model compared to Langdetect baseline. In the meantime, out of the six less popular languages, only two languages (Italian and Slovakian) show an improvement. This might mean a correlation between the performance and the number of training data. However, there is a clear exception. Main model achieves an F1-score of 0.347 for Slovakian, which both AO model and TO model completely failed.

For some languages, Main model achieved a lower performance than a single modality model, leaving room for further improvement. For example, Main model is outperformed by AO model for Korean and TO model for French and Italian.




Figure 6: The comparison of metrics of Main model and models with various modality dropout strategies when there is no missing modalities in the data. ADr-Main, TDr-Main, and ATDr-Main indicate models with the same architecture but with modality dropouts applied on audio input, text input, and both of the inputs, respectively.

5.4 Modality Dropout

In this section, we discuss the benefits and effects of applying modality dropout to Main model. We test three modality dropout strategies: applying modality dropout to audio input only (ADr-Main model), text input only (TDr-Main model), and both of the inputs (ATDr-Main model). The dropout rate is fixed to 0.2 for every model.⁵ In ATDr-Main model, the two dropouts work independently. This means that during training, stochastically, 4% of training items would have zero values for both of the inputs where backpropagation is still applied to update the model. This leads to strengthen the model to predict the distribution of the training data and may not be ideal, but we did not observe any critical issue in practice.

Since there is no language-specific pattern, we present the averaged metrics only in this section.

5.4.1 Case 1: Complete Modality – Do modality dropouts have any negative affects?

It is unusual to use modality dropout in music information retrieval. We first investigate to ensure that adopting it does not harm the normal use-cases, i.e., when there is no missing modality.

As presented in Figure 6, all the models with modality dropouts - ADr-Main model, TDr-Main model, and ATDr-Main model achieve comparable or even outperforming performances over Main model although they might not be statistically significant. The effects of Modality dropout are better reflected on weighted-average scores than macro-average ones. That is because weighted-averaged scores are more linearly related to the empirical loss that our models are trained to minimize. As shown, modality dropouts only improved those scores. To summarize, adopting modality dropouts does not harm the normal use-cases.

5.4.2 Case 2: Missing Audio Input

Figure 7 presents the performances of Main model, ADr-Main model, and ATDr-Main model when audio inputs are missing, i.e., there is only text input.

⁵ 0.2 was chosen assuming a small portion (for example, 20%) of tracks would have missing modality in the real use-cases.

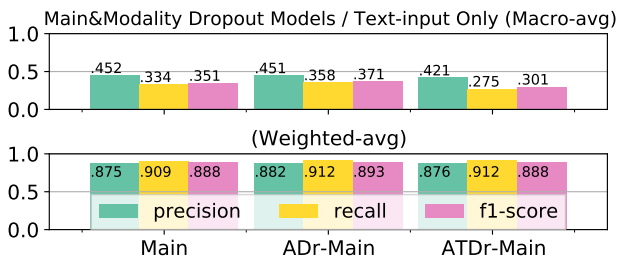


Figure 7: The comparison of metrics of Main model and models with various modality dropout strategies (see 6 for details) when there is missing audio modality. Note that as in Figure 3 (top), TO model, a dedicated text-only model achieved precision / recall / F1-score of .526 / .415 / .422 (macro averaging) and .896 / .914 / .900 (weighted averaging), respectively.




Figure 8: The comparison of metrics of Main model and models with various modality dropout strategies (see 6 for details) when there is missing text modality. Note that as in Figure 3 (bottom), AO model, a dedicated audio-only model achieved precision / recall / F1-score of .387 / .248 / .275 (macro averaging) and .852 / .884 / .857 (weighted averaging), respectively.

First, surprisingly, Main model - one that is trained without any modality dropout - performs comparably with ADr-Main model and ATDr-Main model, which are trained explicitly to be able to deal with missing audio input. This is an unexpected behavior since there is no reason the model should learn to ignore a zero audio input (silence) and make a correct prediction solely based on the text information – but, seems like it does. We find it difficult to explain it and leave it as a future work.

Second, it is worth comparing these results with that of TO model (Figure 3, top), a model that is designed and trained to work with text input only. In all the metrics and averaging strategies, TO model outperforms all the three models. For F1-score, even with the best model (ADr-Main), there is a difference of 0.051 (macro averaging) or 0.007 (weighted averaging). This indicates that despite versatility, a model with an audio modality dropout may not completely replace a dedicated text-only model, especially if missing audio inputs are highly likely.

5.4.3 Case 3: Missing Text Input

Figure 8 presents the performance of Main model, TDr-Main model, and ATDr-Main model when text input is missing, i.e., there is only audio input.

First, there are noticeable improvements by apply-

ing modality dropouts - TDr-Main model and ATDr-Main model outperformed Main model for the most of the metrics. ATDr-Main model achieved +0.086 and +0.101 higher F1-scores than Main model does. This result supports a potential real-world use-case of serving a single SLID model where metadata may or may not be available.

Second, when compared to AO model (Figure 3, top), all the three models in Figure 8 are outperformed. This means, similar to the conclusion of Section 5.4.2, a model with a text modality dropout may not serve as a perfect alternative and whether to apply text modality dropout (as opposed to train two different models, AO model and Main model) would be a practical choice: the decision would be based on the ratio of missing-text inputs and the costs of training and maintaining one vs. two models.

6. CONCLUSION

In this paper, we presented LRID-Net, a deep learning model for singing language identification (SLID) that takes advantage of multimodal data. LRID-Net takes an audio input as well as a text input that combines track title, album name, and artist name. We also propose modality dropout in MIR task, which is designed to let a single model be used with varying input availability. In the experiment, we showed that i) multimodal input improves the performance, ii) a language probability vector of metadata is an effective representation for SLID, iii) modality dropouts do not harm the performance when both of the input modalities exist, and iv) modality dropouts make a model robust with missing input to some extent.

Our research has several limitations. There are some behaviors that we could not provide a satisfying explanation about. Although being useful to some extent, the modality dropout did not completely fulfill the need of building multiple models to cope with missing modalities. Due to the already complicated experiment configuration, we did not opt for balancing the languages, which would be a necessary step to build a more practical language classifier.

There are many research questions to be answered in SLID. A data-driven SLID model might learn some non-linguistic features that are correlated to language labels, and identifying those mechanisms would lead to building more robust SLID models. One approach to demystifying their behavior is to use source separation techniques and observe how much a model is relying on vocal parts vs accompaniments. Source separation techniques also could lead to a better performing model because separated signals would provide a disentangled input representation that may be useful for the task. Finally, another highly related and interesting task is lyric transcription, which can be combined with SLID, where a mutual benefit is anticipated.

7. ACKNOWLEDGEMENT

We thank Jeong Choi, Minz Won, Jordan Smith, Janne Spijkervet, Gianluca Micchi, and Zhihao Ouyang for their helpful reviews and discussion on this paper.

8. REFERENCES

- [1] J. H. Lee and J. S. Downie, "Survey of music information needs, uses, and seeking behaviours: preliminary findings." in *The 5th International Society of Music Information Retrieval Conference (ISMIR)*, vol. 2004. Citeseer, 2004, p. 5th.
- [2] M. Sugiyama, "Automatic language recognition using acoustic features," in *[Proceedings] ICASSP 91: 1991 International Conference on Acoustics, Speech, and Signal Processing*. IEEE, 1991, pp. 813–816.
- [3] L. F. Lamel and J.-L. Gauvain, "Language identification using phone-based acoustic likelihoods," in *Proceedings of ICASSP'94. IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. 1. IEEE, 1994, pp. 1–293.
- [4] W.-H. Tsai, H.-M. Wang *et al.*, "Towards automatic identification of singing language in popular music recordings." in *International Society of Music Information Retrieval (ISMIR)*. Citeseer, 2004.
- [5] W.-H. Tsai and H.-M. Wang, "Automatic identification of the sung language in popular music recordings," *Journal of New Music Research*, vol. 36, no. 2, pp. 105–114, 2007.
- [6] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front-end factor analysis for speaker verification," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 4, pp. 788–798, 2010.
- [7] A. M. Kruspe and I. Fraunhofer, "Improving singing language identification through i-vector extraction." in *DAFx*, 2014, pp. 227–233.
- [8] A. M. Kruspe, J. Abesser, and C. Dittmar, "A gmm approach to singing language identification," in *Audio Engineering Society Conference: 53rd International Conference: Semantic Audio*. Audio Engineering Society, 2014.
- [9] J. Schwenninger, R. Brueckner, D. Willett, and M. E. Hennecke, "Language identification in vocal music." in *International Society of Music Information Retrieval (ISMIR)*, 2006, pp. 377–379.
- [10] D. DEL CASTILLO, "End-to-end learning for singing-language identification," 2020, master's thesis, KTH, School of Electrical Engineering and Computer Science (EECS).
- [11] K. Drossos, S. I. Mimilakis, D. Serdyuk, G. Schuller, T. Virtanen, and Y. Bengio, "Mad twinnet: Masker-denoiser architecture with twin networks for monaural sound source separation," in *2018 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2018, pp. 1–8.
- [12] V. Chandrasekhar, M. E. Sargin, and D. A. Ross, "Automatic language identification in music videos with low level audio and visual features," in *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2011, pp. 5724–5727.
- [13] L. Roxbergh, "Language classification of music using metadata," 2019, master's thesis, Uppsala Universitet.
- [14] I. A. P. Santana, F. Pinhelli, J. Donini, L. Catharin, R. B. Mangolin, V. D. Feltrim, M. A. Domingues *et al.*, "Music4all: A new music database and its applications," in *2020 International Conference on Systems, Signals and Image Processing (IWSSIP)*. IEEE, 2020, pp. 399–404.
- [15] T. Bertin-Mahieux, D. P. Ellis, B. Whitman, and P. Lamere, "The million song dataset," 2011.
- [16] K. Choi, G. Fazekas, and M. Sandler, "Automatic tagging using deep convolutional neural networks," *The 17th International Society of Music Information Retrieval Conference, New York, USA*, 2016.
- [17] K. Choi, G. Fazekas, K. Cho, and M. Sandler, "The effects of noisy labels on deep convolutional neural networks for music tagging," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 2, no. 2, pp. 139–149, 2018.
- [18] N. Shuyo, "Language detection library for java," 2010. [Online]. Available: <http://code.google.com/p/language-detection/>
- [19] K. Choi, D. Joo, and J. Kim, "Kapre: On-gpu audio preprocessing layers for a quick implementation of deep neural network models with keras," in *Machine Learning for Music Discovery Workshop at 34th International Conference on Machine Learning*. ICML, 2017.
- [20] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [21] M. Won, A. Ferraro, D. Bogdanov, and X. Serra, "Evaluation of cnn-based automatic music tagging models," *Proceedings of 17th Sound and Music Computing (SMC)*, 2020.
- [22] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *International conference on machine learning*. PMLR, 2015, pp. 448–456.
- [23] N. Neverova, C. Wolf, G. Taylor, and F. Nebout, "Mod-drop: adaptive multi-modal gesture recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 8, pp. 1692–1706, 2015.
- [24] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *3rd International Conference for Learning Representations, San Diego*, 2015.