

# Innovations in Cover Song Detection: A Lyrics-Based Approach

Maximilian Balluff<sup>1</sup>, Peter Mandl<sup>1</sup>, Christian Wolff<sup>2</sup>,

<sup>1</sup>Hochschule Muenchen University of Applied Science, <sup>2</sup>University of Regensburg,

Correspondence: [maximilian.balluff@hm.edu](mailto:maximilian.balluff@hm.edu)

## Abstract

Cover songs are alternate versions of a song by a different artist. Long being a vital part of the music industry, cover songs significantly influence music culture and are commonly heard in public venues. The rise of online music platforms has further increased their prevalence, often as background music or video soundtracks. While current automatic identification methods serve adequately for original songs, they are less effective with cover songs, primarily because cover versions often significantly deviate from the original compositions. In this paper, we propose a novel method for cover song detection that utilizes the lyrics of a song. We introduce a new dataset for cover songs and their corresponding originals. The dataset contains 5078 cover songs and 2828 original songs. In contrast to other cover song datasets, it contains the annotated lyrics for the original song and the cover song. We evaluate our method on this dataset and compare it with multiple baseline approaches. Our results show that our method outperforms the baseline approaches.

## 1 Introduction

Cover versions, or cover songs, represent reinterpretations of original compositions by different artists, holding a significant place in the music industry and influencing musical culture. These renditions are frequently featured in public venues such as bars, clubs, festivals, and restaurants. The rise of online music platforms has further amplified their ubiquity, often serving as background music or soundtracks for videos (*Vervielfältigungsrechte, 2021*).

Identifying and monetizing cover songs poses a significant challenge for rights holders, leading to an increasing demand for efficient methods to detect cover songs on a large scale. Composers and performers, often members of collecting societies, rely on accurate song identification for royalty distribution. While existing automatic identification

methods effectively handle original songs, detecting cover songs remains challenging due to the substantial variations introduced by different artists in their renditions (*Sha, 2023*).

## 2 Related Work

Cover Song Detection (CSD) plays a pivotal role in Music Information Retrieval, especially in recent years. Existing approaches are predominantly centered around audio analysis. For instance, Tsai et al. (*Tsai et al., 2005*) focused on tempo, transposition, and accompaniment to compare original songs with their covers. Serrà introduced a method using audio features to compute a similarity matrix between original and cover songs (*Serrà Julià, 2007*). In another study (*Du et al., 2021*), advanced song embeddings were generated using a convolutional neural network and utilized for similarity computation. These methods commonly extract audio features, ranging from high-level attributes like tempo and beats per minute (*Tsai et al., 2005*) to low-level features like Harmonic Pitch Class Profile (*Serrà Julià, 2007*) and Constant-Q-Transform (*Du et al., 2021*). Some approaches combine both, incorporating features like the mean and standard deviation of Mel-Frequency Cepstral Coefficient (*Yu et al., 2008*).

In contrast, only a few methods leverage song lyrics for CSD. Correya et al. employed a unique approach using the lyrics and title of songs through a bag-of-words algorithm (*Correya et al., 2018*). Similarly, Vaglio et al. proposed a method combining lyrics and audio for cover song detection (*Vaglio et al., 2021*). They processed the audio in two branches—one for lyrics and the other for audio. After detecting singing voice, lyrics were transcribed, and string matching identified corresponding cover versions. The results from both branches were then combined for the final output.

### 3 Idea and Dataset

Our central proposition asserts that cover song lyrics exhibit notable similarities with their original counterparts. Given the pivotal role of lyrics in music, especially in genres like Pop, Hip Hop/Rap, Rock/Hardrock/Heavy Metal, and Kids Music, where over 70% of sales are concentrated (e. V., 2021), investigating this hypothesis becomes imperative.

To substantiate our claim, we introduce a novel dataset comprising 5078 cover songs and 2828 original songs. Collected from genius.com between June and September 2023, this dataset stands out by providing annotated lyrics for both the original and cover songs. Language analysis, conducted using a transformer for language detection (Conneau et al., 2020), reveals that approximately 85% of the original songs are in English, followed by Spanish (4.7%) and Portuguese (2.3%).

To assess lyrical similarity, we employ Levenshtein distance (Levenshtein, 1996) and word error rate (WER) (Klakov and Peters, 2002). The mean Levenshtein distance between a cover song and its original 302.62 is notably lower than the mean (689.25) and minimum (504.76) distances to other songs; for more details, compare figure 1. Surprisingly, 23.6% of non-matching songs exhibit a lower distance than the actual original, with 301 instances where the detected language of the cover and original differs. Further analysis reveals instances of low distances in non-matching songs, often attributed to poor annotations, including incomplete cover lyrics or variations in annotation styles (e.g., filler words like "Uh, uh, uh..." or paragraph descriptions like [Verse], [Refrain]). The average WER distance between a cover song and its original, denoted as 0.54, is significantly less than the average (1.33) and minimum (0.87) distances observed with other songs. These findings underscore the challenges posed by these differences for classical distance metrics, necessitating thorough preprocessing for accurate analysis.

The dataset will be published in the future. Currently, we are working on increasing the size of the dataset and improving the quality. If possible we will merge our dataset with other datasets like Second Hand Songs (SHS) (Bertin-Mahieux et al., 2011).




Figure 1: Levenshtein distance between (upper) and WER (lower) between cover and original songs, and between cover and other songs

#### 3.1 Baseline

As we are the first to utilize our new dataset, our primary objective is to establish multiple baseline algorithms for comparison purposes. We utilize the string similarity distances, namely Levenshtein distance ( $L_D$ ) and WER ( $WER$ ), along with a Bag-of-Words algorithm combined with Cosine distance ( $BOW$ ), to calculate the similarity between lyrics. We opted for this set of distances due to their relatively low complexity, and we benefit from existing implementations for Levenshtein and WER (Bachmann, 2023; pyt, 2023).

#### 3.2 Proposed method

Since the initial introduction of transformer neural networks (Vaswani et al., 2017), these networks have become the standard for text preprocessing. We utilize a transformer as a feature extractor for the lyrics, which generates an embedding. Lyrics with similar content should yield similar embeddings. The embeddings for each original song are stored in a database. For each cover song, we compute the similarity between the cover song's embedding and the embeddings of the original songs. We order the songs by their similarity, the original song with the highest similarity score is predicted to be the original song for the cover version.

##### 3.2.1 Model

In our approach, we use a pre-trained model for sequence similarity (Müller et al., 2022), which is




Figure 2: Model architecture

based on XLM-RoBERTa (Conneau et al., 2020) and is pre-trained on the SNLI (Bowman et al., 2015), MNLI (Williams et al., 2018), ANLI (Nie et al., 2020) and XNLI (Conneau et al., 2018) datasets. On top of the base model, we add two fully connected and an embedding layer. The embedding layer serves the purpose of transforming the sequence of features into a fixed-length floating-point embedding vector, providing a condensed representation of the lyrics that enables subsequent similarity calculations. The model architecture is shown in figure 2. We did not further optimize the model architecture. The dimensions of the layers are based on our experience in previous text matching projects.

Various methods can be employed to train embeddings, and a common approach is to use Siamese neural networks (Chicco, 2021). These networks consist of two identical neural networks that share the same weights. The inputs for these networks are the two samples that need to be compared, with the objective of minimizing the distance between similar inputs and maximizing the distance between dissimilar ones.

In our approach, we adopt the triplet loss method, originally introduced for face recognition (Chechik et al., 2009). The goal is to minimize the distance between the anchor and the positive sample while maximizing the distance between the anchor and the negative sample. Achieving this involves sampling triplets, which consist of an anchor, a positive sample, and a negative sample (Chechik et al., 2009). To address the challenge of the relatively high cost of finding good triplets, we employ the online mining strategy (Schroff et al., 2015). Instead of mining triplets before passing them to the model, the triplets are mined within the training batch. This is accomplished by passing a batch of samples to the model, calculating the embeddings and distances for all embeddings in the batch, and then searching for each sample used as an anchor

for other positive and negative samples within the batch.

### 3.2.2 Evaluation

The evaluation of our approach relies on three widely-applied metrics in the field of Cover Song Detection (CSD): mean average precision (mAP), mean rank (MR), and precision at one (P@1). These metrics have been extensively used in the evaluation of various CSD approaches (Du et al., 2021, 2022; Chicco, 2021). mAP calculates the average precision of a ranked list of documents, where precision is defined as the ratio of relevant documents among the retrieved documents (Christopher et al., 2008). MR provides the position of the first relevant document in the ranked list of documents, with lower values indicating better performance (ideally one) (Osmalskyj, 2017). P@1 represents the precision at the first rank of the ranked list of documents (Xu et al., 2018).

Given the limited size of our dataset, we chose to implement k-fold cross-validation with five folds, as recommended by (Dat, 2011). In order to maintain consistency and comparability across all experiments, the same folds were utilized. Each fold encompassed several cover lyrics and their corresponding originals. The fold was further divided into training (64%), validation (16%), and testing (20%) sets. The validation set is necessary to prevent our model from overfitting the training data, using EarlyStopping and ReduceLearningRate callbacks that monitor the validation loss.

It is crucial to note that no original version is duplicated among these sets, ensuring that each set contains a unique set of original songs. This resulted in varying sizes for each fold, due to the uneven distribution of the number of covers per original. Table 1 provides an overview of the number of unique lyrics (original and cover) for each fold.

To evaluate the model’s performance, we utilized the test set. For each lyric, whether it be a cover or an original, we generated embeddings. The embeddings from the original lyrics formed our database. Subsequently, for each cover embedding, we calculated the Manhattan distance (Royer, 2001) to all original embeddings, sorting them based on their distances. Our metrics were then computed using this ranking. Model training was conducted on an Nvidia T4 GPU, utilizing a n1-standard-8 instance on Google Cloud. The training process spanned approximately 22 hours for all five folds, averag-

	Training	Validation	Test
<b>Fold 1</b>	5058	1313	1535
<b>Fold 2</b>	4975	1397	1534
<b>Fold 3</b>	5069	1224	1613
<b>Fold 4</b>	4953	1334	1619
<b>Fold 5</b>	5026	1275	1605

Table 1: Unique lyrics (original and cover) for each fold

	mAP	MR	P@1
<b>L<sub>D</sub></b>	81.65 ± 0.74%	82.57 ± 4.74	80.23 ± 0.54%
<b>WER</b>	82.91 ± 1.27%	75.01 ± 8.43	81.52 ± 0.97%
<b>BOW</b>	85.74 ± 0.67%	46.29 ± 1.80	<b>83.65</b> <b>± 0.71%</b>
<b>Triplet Model</b>	<b>87.17</b> <b>± 1.50%</b>	<b>18.51</b> <b>± 2.35</b>	83.57 ± 1.80%

Table 2: Average results on the 5-fold cross-validation

ing around 19 epochs. Notably, our model boasts 278 million parameters, of which 831 thousand are trainable.

Our baseline approaches do not require training, and therefore, we only use the test set for evaluation. We compute the distance between the cover song and all original songs. The original song with the smallest distance is then predicted as the original for the cover song.

### 3.2.3 Results

The results of our experiments are summarized in Table 2. The Bag-of-Words (*BOW*) approach emerges as the best baseline, yielding a mAP of 85.74%, a MR of 46.29, and Precision@1 of 83.65%. The triplet model outperforms the baseline approaches in all metrics except precision, achieving a mAP of 87.17%, a MR of 18.51, and Precision@1 of 83.57%. Notably, all methods yield approximately the same mAP, while the MR varies significantly. Remarkably, the triplet model, with its substantially lower rank, outperforms the *BOW* model by more than a factor of two.

Although we did not initially plan to evaluate the computation time of the approaches, it is worth noting that WER distance is significantly slower. It took over 8 hours to compute the distance for all splits, whereas for all other approaches, it was a matter of minutes.

## 4 Limitations and Future Work

While our approach shows promise, it is not without limitations. Firstly, direct comparisons with other CSD approaches that focus on audio features pose challenges, given our method’s exclusive operation on our introduced dataset. A meaningful comparison would necessitate additional steps, including downloading and extracting audio files from YouTube, conducting speech-to-text transcriptions, and comparing lyrics. Importantly, a comparison of our approaches to methods that use datasets like SHS (Bertin-Mahieux et al., 2011) requires evaluating how well current speech-to-text algorithms handle singing voice, potentially involving training models like Whisper (Radford et al., 2022) for this task. This leads to the second limitation, lyrics extracted from the audio may differ from the annotated lyrics, introducing potential errors, or missing annotation. As already mentioned it is necessary to evaluate our approaches on extract lyrics. Lastly, the context window of our model, while suitable for average-sized lyrics, may prove insufficient for larger lyrics, potentially leading to information loss and impacting cover song detection accuracy.

In future work, addressing these limitations is crucial. Expanding the dataset size could enhance the model’s performance. Furthermore, we intend to extend training to the entire model, encompassing not only the fully connected and embedding layers. Experimenting with direct lyric extraction from audio, especially from datasets like SHS, would facilitate meaningful comparisons with other CSD methods. To tackle the context window limitation, research efforts could focus on methods for handling larger lyrics, such as expanding the context window or exploring alternative solutions like splitting lyrics into smaller segments for more effective processing. It might also be possible to partly merge our dataset with other datasets like SHS to increase the size of the dataset.

## 5 Responsible NLP

In light of the current state of our methodology, it has been conclusively ascertained that the implementation of a large-scale CSD system in a professional environment is currently unfeasible. The metrics provided herein illustrate specific circumstances wherein our methodology fails to produce accurate results, thus introducing a potential risk of incorrect matches. Such inaccuracies could potentially disrupt the fair distribution of royal-

ties. Nevertheless, there are potential solutions that could mitigate this risk. For instance, imposing a distance limit and only considering automated matches within close proximity could be a viable strategy. We acknowledge that transformer-based models inherently possess biases that may facilitate more effective matching of certain words or phrases over others. Although there is no concrete evidence at present to suggest bias in our model, we unequivocally pledge to comprehensively investigate this issue in our future research endeavors. With respect to the reproducibility of our approach’s results, we do not have concerns. Given that our model does not generate text and as such does not possess a hyperparameter for output generation, our results are reproducible. For the same given lyrics, the same embedding vector will be produced, thereby yielding the identical original lyrics. In the development of our model, we utilized GitHub Copilot. It is important to note that this tool did not exert any influence on our research, and we carefully reviewed the generated code. Additionally, we employed GPT-4 to check for spelling and grammatical inaccuracies present in our manuscript and correct them.

## References

2011. **Front Matter**. In Ian H. Witten, Eibe Frank, and Mark A. Hall, editors, *Data Mining: Practical Machine Learning Tools and Techniques (Third Edition)*, The Morgan Kaufmann Series in Data Management Systems, pages i–iii. Morgan Kaufmann.
2023. **Shazam**.
2023. *WordErrorRate - Torch Eval*. PyTorch.
- Max Bachmann. 2023. *Levenshtein*. <https://github.com/maxbachmann/Levenshtein>.
- Thierry Bertin-Mahieux, Daniel P.W. Ellis, Brian Whitman, and Paul Lamere. 2011. The million song dataset. In *Proceedings of the 12th International Conference on Music Information Retrieval (ISMIR 2011)*.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. **A large annotated corpus for learning natural language inference**. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642. Association for Computational Linguistics.
- Gal Chechik, Varun Sharma, Uri Shalit, and Samy Bengio. 2009. **Large Scale Online Learning of Image Similarity through Ranking**. In Helder Araujo, Ana Maria Mendonça, Armando J. Pinho, and María Inés Torres, editors, *Pattern Recognition and Image Analysis*, volume 5524, pages 11–14. Springer Berlin Heidelberg.
- Davide Chicco. 2021. **Siamese Neural Networks: An Overview**. In Hugh Cartwright, editor, *Artificial Neural Networks*, Methods in Molecular Biology, pages 73–94. Springer US.
- Manning Christopher, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. **Unsupervised Cross-lingual Representation Learning at Scale**.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel R. Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. **Xnli: Evaluating cross-lingual sentence representations**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Albin Andrew Correya, Romain Hennequin, and Mickaël Arcos. 2018. **Large-Scale Cover Song Detection in Digital Music Libraries Using Metadata, Lyrics and Audio Features**.
- Xingjian Du, Ke Chen, Zijie Wang, Bilei Zhu, and Zejun Ma. 2022. **Bytecover2: Towards Dimensionality Reduction of Latent Embedding for Efficient Cover Song Identification**. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 616–620. IEEE.
- Xingjian Du, Zhesong Yu, Bilei Zhu, Xiaoou Chen, and Zejun Ma. 2021. **Bytecover: Cover Song Identification Via Multi-Loss Training**. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 551–555.
- Bundesverband Musikindustrie e. V. 2021. **Musikindustrie in Zahlen im Überblick: BVMI - 2021**.
- Dietrich Klakow and Jochen Peters. 2002. **Testing the correlation of word error rate and perplexity**. *Speech Communication*, 38(1):19–28.
- V. I. Levenshtein. 1996. **Binary Codes Capable of Correcting Deletions, Insertions and Reversals**. 10:707.
- Thomas Müller, Guillermo Pérez-Torró, and Marc Franco-Salvador. 2022. **Few-Shot Learning with Siamese Networks and Label Tuning**. In *ACL (to appear)*.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. **Adversarial NLI: A new benchmark for natural language understanding**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.

- Julien Osmalskyj. 2017. [A Combining Approach to Cover Song Identification](#).
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. [Robust speech recognition via large-scale weak supervision](#).
- Christian Royer. 2001. *Simultane Optimierung von Produktionsstandorten, Produktionsmengen und Distributionsgebieten*. Herbert Utz Verlag.
- Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. [FaceNet: A Unified Embedding for Face Recognition and Clustering](#). In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 815–823.
- Joan Serrà Julià. 2007. Music similarity based on sequences of descriptors tonal features applied to audio cover song identification.
- Wei-Ho Tsai, Hung-Ming Yu, and Hsin-min Wang. 2005. [A Query-By-Example Technique for Retrieving Popular Songs with Similar Melodies](#). In *ISMIR 2005, 6th International Conference on Music Information Retrieval*, page 190.
- Andrea Vaglio, Romain Hennequin, Manuel Moussallam, and Gael Richard. 2021. [THE WORDS REMAIN THE SAME: COVER DETECTION WITH LYRICS TRANSCRIPTION](#). In *22nd International Society for Music Information Retrieval Conference ISMIR 2021*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- GEMA-Gesellschaft für musikalische Aufführungs- und mechanische Vervielfältigungsrechte. 2021. [Diskotheken-monitoring](#).
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics.
- Xiaoshuo Xu, Xiaou Chen, and Deshun Yang. 2018. [Key-Invariant Convolutional Neural Network Toward Efficient Cover Song Identification](#). In *2018 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6.
- Yi Yu, J. Stephen Downie, Fabian Moerchen, Lei Chen, and Kazuki Joe. 2008. [Using Exact Locality Sensitive Mapping to Group and Detect Audio-Based Cover Songs](#). In *2008 Tenth IEEE International Symposium on Multimedia*, pages 302–309.