

Leveraging User-Generated Metadata of Online Videos for Cover Song Identification

Simon Hachmeier and Robert Jäschke

School of Library and Information Science

Humboldt-Universität zu Berlin

{simon.hachmeier, robert.jaeschke}@hu-berlin.de

Abstract

YouTube is a rich source of cover songs. Since the platform itself is organized in terms of videos rather than songs, the retrieval of covers is not trivial. The field of cover song identification addresses this problem and provides approaches that usually rely on audio content. However, including the user-generated video metadata available on YouTube promises improved identification results. In this paper, we propose a multi-modal approach for cover song identification on online video platforms. We combine the entity resolution models with audio-based approaches using a ranking model. Our findings implicate that leveraging user-generated metadata can stabilize cover song identification performance on YouTube.

1 Introduction

Music is a popular content category on YouTube (Montero and Mora-Fernandez, 2020). Uploaders share music in a variety of contexts, ranging from amateur covers to mashups (Airoldi et al., 2016; Liikkanen and Salovaara, 2015). Since YouTube is not organized in terms of songs but rather in terms of online videos,¹ finding cover versions of songs is a non-trivial retrieval task.

Driven by applications such as copyright infringement detection, cover song identification (CSI) deals with the retrieval of covers. The key challenge of CSI is to compare songs based on properties which can indicate their association (e.g., melody, lyrics) while discarding irrelevant information (e.g., timbre). Consequently, current research efforts are mainly audio-based (Du et al., 2023; Liu et al., 2023; Hu et al., 2022; Yu et al., 2020) with limited consideration of non-audio features such as lyrics (Abrassart and Doras, 2022). However, the utilization of user-generated metadata of online videos like in similar tasks (Agrawal and Sureka;

Smith et al., 2017), has not yet been considered in CSI.

We propose to model CSI on online video platforms (OCSI) as a multi-modal problem, based on the hypothesis that uploaders tend to describe their videos using attributes of songs (e.g., song title, performer name) to make them easily findable. In this work, we propose multi-modal ensembles combining entity resolution models with audio-based CSI models in a late-fusion fashion using the ranking model LambdaMART (Wu et al., 2010). We compare the performance of the proposed ensembles with the performance of the CSI models. Further, we study the robustness of ER approaches in difficult cases such as song title variations or hard negatives on YouTube and provide our code and results.²

2 Multi-Modal Online Cover Song Identification

The items in our task are online videos. The goal of OCSI is the retrieval of items associated with the same musical work as a query item. In traditional CSI, only representations concerning musical content (e.g., audio and lyrics) are considered. We additionally leverage user-generated metadata. In Figure 1 we show an example of a query-candidate pair.

Each item is represented by attributes derived from its audio data and attributes from its user-generated metadata (video title, channel name, video description, and a set of keywords). For simplicity, we assume that each item contains only one song. A song has the attributes song title, performer name, and a work identifier. Songs which are associated with the same musical work are considered relevant in the retrieval scenario when one song is used as a query to retrieve the other song.

We model the task of OCSI as a multi-modal

¹Except for YouTube’s streaming service *YouTube Music*.

²https://github.com/progsi/er_csi

	Query Item	Candidate Item
Data Source:	Secondhandsongs	
	Song Title: Yesterday Performer: Jan & Dean	Song Title: Yesterday Performer: Mary Wells
Data Source:	YouTube	
	Video Title: Yesterday Channel: Jan & Dean Topic Descr: Provided to ...	Video Title: Mary Wells - Yesterday Channel: Franklin Pierce Descr:
Input		
Simple	Yesterday	Mary Wells - Yesterday
Rich	[COL] title [VAL] Yesterday [COL] performer [VAL] Jan & Dean [COL] video [VAL] Yesterday [COL] channel [VAL] Jan & Dean - Topic [COL] description [VAL] Provided to YouTube by Universal Music Group	[COL] title [VAL] [MASK] [COL] performer [VAL] [MASK] [COL] video [VAL] Mary Wells - Yesterday [COL] channel [VAL] Franklin Pierce [COL] description [VAL]

Figure 1: Example of the input of items of the work “Yesterday” written by John Lennon and Paul McCartney. Colors in the box frames and text indicate the data source: blue stands for Secondhandsongs and red for YouTube.

problem involving metadata and audio representations. Our overall system will consist of modules to compute the pairwise similarities for each of both modalities. Then, a ranking model combines both outputs to compute an overall rank. In the following, we explain our entity-resolution (ER) methods used to model similarities in the metadata domain.

2.1 Entity Resolution

Fuzzy Matching We use the token ratio function from rapidfuzz (Bachmann, 2021), which turned out to be the best performing fuzzy matcher for the task in a preliminary experiment. For a pair of strings, the function returns the maximum between their normalized Indel similarity and the token set ratio. The former is the minimum number of insertions and deletions to convert one string into the other. The latter is the number of tokens in the shorter string contained in the longer string divided by the number of tokens in the longer string. We validated the use of the token ratio by the MAP on the validation dataset (cf. Section 3). We tested matching song title and performer concatenated and using solely the song title of the query item. For the candidate item, we experimented with only the video title and with the latter combined with the other attributes concatenated by space. The best configuration was simply matching the song title to the video title (cf. *Simple* input in Figure 1). We also experimented with the snowball stemmer from NLTK (Bird et al., 2009) for all attributes but it did not improve the results.

S-BERT The model S-BERT (Reimers and Gurevych, 2019) addresses the problem of quadratic complexity of language models (LMs) like BERT (Devlin et al., 2018) when processing pairs of sentences. Hence, it was used as promising method for ER (Li et al., 2021a; Paganelli et al., 2023). The model learns to encode sentences into embedding vectors which can be compared using similarity measures such as Cosine similarity.

We fine-tune a multilingual approach of S-BERT (Reimers and Gurevych, 2020) which encodes text sequences into 384-dimensional vectors. Like for fuzzy matching, we use the *Simple* input because it performed better than the other attribute combinations. We apply a similar training procedure like in recent CSI approaches: we use triplet loss with a margin of 0.3 and apply online hard triplet mining (Xuan et al., 2020) where the hardest triplets in the batch of 16 items with 4 random works represented by 4 items are used for training updates. We select the best model after 10 epochs measured in MAP on the validation dataset.

Ditto Other state-of-the-art approaches rely on contextualized embeddings provided by pretrained LMs to predict the matching confidence for a given entity pair. In theory, this can improve the ER task since the context is not only considered for tokens in one entity but across both entities. Thus, we experimented with Ditto (Li et al., 2020), HierGAT (Yao et al., 2022), and r-SupCon (Peeters and Bizer, 2022). We found that Ditto was both – better performing and faster in inference. We therefore select Ditto for our experiments. The model computes a binary matching confidence based on entity pairs encoded by LMs such as RoBERTa (Liu et al., 2019). Due to the quadratic complexity during inference, we use S-BERT as blocker. We adopt the top- k blocking strategy suggested by the authors (Li et al., 2021b) where the top- k most similar candidate items per query item are passed to Ditto and the remaining pairs are predicted by the blocker.

As underlying LM, we chose RoBERTa since it was shown to achieve high performance in the ER task (Li et al., 2020; Peeters and Bizer, 2022; Peeters et al., 2023; Yao et al., 2022) and multilingual variant of BERT (mBERT) (Devlin et al., 2018). First, we experimented with the same attribute combinations we use for fuzzy matching and S-BERT. We observed that Ditto works better

when items are represented with the same attributes on the query and candidate side. Hence, we use a *rich input* (cf. Figure 1). As inputs for the LMs, we concatenate the names and values of attributes of the query and candidate item: Each attribute is represented by a [COL] token followed by its name (e.g., *title* for the song title) and a [VAL] token followed by its value (e.g., *Yesterday*). Since the actual song attributes of the items on the candidate side are not known, we mask the respective tokens of those using a [MASK] token. We fine-tune with a batch size of 32 and a learning rate of 1e-05 and a sequence length of 256 tokens. We select the best performing model after 15 epochs. We fine-tune Ditto with a dataset of pairs (cf. Section 3) and measure the performance on the validation dataset in F1.

2.2 Combining ER with CSI

We form multi-modal ensembles each combining one fine-tuned ER model with a trained CSI model (ER-CSI ensembles). We use two pre-trained CSI models: CQNet (Yu et al., 2020) and CoverHunter (Liu et al., 2023). Both models encode items into vectors and represent musical similarity using Cosine similarity. The former uses convolutional neural networks to learn 300-dimensional vector representations. CoverHunter uses conformer neural networks (Gulati et al., 2020) and an attention mechanism for temporal pooling (Okabe et al., 2018) to learn 128-dimensional vector representations. For each ER-CSI ensemble we train the ranking model LambdaMART (Wu et al., 2010) using the pairwise similarities as input features. We use (mean average precision) MAP objective function and consider the top 50 feature interactions similar to (Lucchese et al., 2022).

3 Experimental Setup

Our experiments aim to evaluate a) whether ensembles of ER and CSI models outperform CSI models and b) whether ER models are robust against hard negatives and song title variations (e.g., translations of song titles or parodies).³ We report two evaluation metrics suggested by MIREX:⁴ MAP and mean rank of the first relevant item (MR1).

We use subsets of two popular CSI datasets: SHS100K (Xu et al., 2018) and DaTacos (Yesiler

³An example for a parody title is “Bye, Bye Johnny” by The Rattles covering “Johnny B. Goode”

⁴cf. https://www.music-ir.org/mirex/wiki/2021:Audio_Cover_Song_Identification

Dataset	Split	Works	Avg.	FM	DI	SB
V-SHS-F	Train	1,501	1.33	✗	✓	✓
V-SHS-V	Valid	1,827	4.73	✓	✗	✓
V-SHS-P	Valid	1,224	1,63	✗	✓	✗
V-SHS-T	Test	1,679	5.25	✓	(✓)	✓
-Unique	Test	852	2.75	✓	✓	✓
-Noise	Test	12	4.15	✓	✓	✓
V-DaT	Test	2,784	4.92	✓	(✓)	✓

Table 1: Dataset statistics: the number of works (Works) with average number of items (Avg.) and usage for Fuzzy Matching (FM), Ditto (DI), or S-BERT (SB). (✓) denotes partial use (after blocking).

et al., 2019), which are popular in CSI research. Both datasets contain items which are songs represented by metadata attributes, a YouTube identifier, and a work identifier. We only retain items with available videos and denote the resulting datasets as V-SHS and V-DaT respectively (cf. Table 1). We retrieve YouTube metadata for all the videos using *YouTube Search Python* and extract the audio features as described by the authors of CoverHunter (Liu et al., 2023) and CQNet (Yu et al., 2020).

Training and Validation Subsets We fine-tune S-BERT and Ditto and train LambdaMART using a random sample of items from the V-SHS100K training set with 1,000 positive pairs and 6,000 negative pairs similar to datasets by Konda et al. (2016). Similarly, we create a pair-wise validation dataset used to select the best model checkpoint of Ditto. The full overview of datasets used is given in Table 1.

Test Subsets Addressing a), we select the *V-DaT* and *V-SHS-T* which are subsets with available videos of datasets typically used in CSI evaluation. For b), the robustness study, we create additional test datasets. First, we only retain one item per work and song title (dropping multiple items with the same song title per work). This subset based on *V-SHS-T* is denoted by *V-SHS-T-Unique*.

Lastly, we aim to evaluate the ER models’ robustness to hard negatives. On YouTube, this can be expected in cases where either the song title is the same for different works or when the words in the song title are used in a different context (e.g., the song title “Hush” occurring in the sentence “Relaxing Hush Sounds”). We focus on the latter problem which is particularly challenging for song titles with one word, because these words are more likely to occur in different contexts. We create

Song Title	Utterances in generated video title
Yesterday	Yesterday 's Kitchen: Old Recipes
Hush	Relaxing Hush Sounds
Time	Mastering Time Management

Table 2: Examples of one-word song titles and video titles of generated hard negatives with ChatGPT 3.5 in the SHS100K-T-Noise dataset.

a subset of *V-SHS-T* containing only items with one-word song titles. For each of the works in the dataset we instruct ChatGPT 3.5 (OpenAI Inc.) to generate video metadata containing the respective word.⁵ We show some examples in Table 2. The resulting dataset *V-SHS-T-Noise* contains 12 works each with 5 generated video titles (hard negatives) and on average around 4 items of the work.

4 Results

In Table 3 we report the results of experiment a). Generally, we observe a strong improvement in MR1 when comparing ER-CSI ensembles with CSI. Improvements in MAP are evident but smaller for ensembles with CoverHunter. Still, gains up to +9% in MAP and up to -13.45 ranks in MR1 are achieved on V-DaTacos. The ensembles with Fuzzy Matching have relatively small performance gains and do not achieve a higher MAP than CoverHunter.

The highest results in MAP are achieved with S-BERT and Ditto with S-BERT as blocker and RoBERTa as underlying LM. Considering the computational overhead for using Ditto, even with blocking and $k = 100$, makes its use for the task questionable. However, Table 4 shows that combining Ditto with S-BERT can stabilize robustness in some cases. Even though the latter achieves higher MAP on the -Unique subsets, combining it with Ditto (M) yields +5% and +8% in MAP on those. Apparently, this might be due to better rankings of Ditto at the earlier ranks indicated by MR1.

5 Limitations and Conclusion

In this paper, we implemented ER approaches as means to support the task of cover song identification in online videos on the example of the online video platform YouTube. We showed that simple fuzzy matching can partly help to increase model performances. Better results were achieved by using S-BERT. Additionally, using Ditto appears to

⁵The prompt was: *Consider the following a list of words. Generate a meaningful video title for each of these words.*

	V-SHS-T		V-DaT	
	MAP	MR1	MAP	MR1
CQNet	0.71	47.40	0.74	10.74
-Fuzzy	0.75	14.92	0.80	4.16
-S-BERT	0.85	12.14	0.91	3.06
-SB+Ditto (R)	0.85	16.29	0.92	3.03
-SB+Ditto (M)	0.83	20.62	0.90	3.49
CoverHunter	0.92	12.60	0.84	15.71
-Fuzzy	0.90	4.47	0.84	5.57
-S-BERT	0.93	3.58	0.93	3.00
-SB+Ditto (R)	0.93	5.41	0.93	2.26
-SB+Ditto (M)	0.92	7.06	0.91	2.70

Table 3: Experiment a): Performances of ER-CSI ensembles against CSI models. SB+Ditto denotes Ditto with S-BERT as blocker; (R) stands for RoBERTa and (M) for mBERT.

	-Noise		SHS-Uniq.	
	MAP	MR1	MAP	MR1
Fuzzy	0.38	5.03	0.37	189.78
S-BERT	0.53	4.15	0.55	138.04
Ditto (R)	0.43	2.97	0.37	114.94
Ditto (M)	0.49	4.21	0.46	100.26
SB+Ditto (R)	0.56	2.97	0.51	183.22
SB+Ditto (M)	0.44	7.31	0.60	252.31

Table 4: Experiment b): Results of ER approaches on the V-SHS-T-Noise (-Noise), V-SHS-T-Unique (S-Uniq.). Due to the smaller dataset size, we set $k = 10$ for -Noise.

be adequate only in some cases such as song title variations. However, the robustness of ER models is generally harmed by hard negatives which potentially do not refer to music.

Lastly, we outline limitations of this study. Our selected text input structure for S-BERT and fuzzy matching can only detect the song title in the video titles. While this might be sufficient for our utilized datasets, other videos which contain the title only in the description or keywords are not uncovered. Secondly, song titles that are completely different than the reference (e.g., parodies or medleys) cannot be detected by ER models. Hence, we see these models as supporting tool in OCSI rather than independent approaches. In future research, we aim to leverage more recent large LMs which are starting to get used for ER (Peeters and Bizer, 2024).

References

- Mathilde Abrassart and Guillaume Doras. 2022. And what if two musical versions don't share melody, harmony, rhythm, or lyrics? In *International Society for Music Information Retrieval Conference*.
- Swati Agrawal and Ashish Sureka. [Copyright Infringement Detection of Music Videos on YouTube by Mining Video and Uploader Meta-data](#). In *Big Data Analytics*, Lecture Notes in Computer Science, pages 48–67. Springer International Publishing.
- Massimo Airoidi, Davide Beraldo, and Alessandro Gandini. 2016. [Follow the algorithm: An exploratory investigation of music on youtube](#). *Poetics*, 57:1–13.
- Max Bachmann. 2021. maxbachmann/rapidfuzz: Release 1.8. 0.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. " O'Reilly Media, Inc."
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- Xingjian Du, Zijie Wang, Xia Liang, Huidong Liang, Bilei Zhu, and Zejun Ma. 2023. [Bytecover3: Accurate cover song identification on short queries](#). In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5.
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, et al. 2020. Conformer: Convolution-augmented transformer for speech recognition. *arXiv preprint arXiv:2005.08100*.
- Shichao Hu, Bin Zhang, Jinhong Lu, Yiliang Jiang, Wucheng Wang, Lingcheng Kong, Weifeng Zhao, and Tao Jiang. 2022. [WideResNet with Joint Representation Learning and Data Augmentation for Cover Song Identification](#). In *Proc. Interspeech 2022*, pages 4187–4191.
- Pradap Konda, Sanjib Das, Paul Suganthan G. C., An-Hai Doan, Adel Ardalan, Jeffrey R. Ballard, Han Li, Fatemah Panahi, Haojun Zhang, Jeff Naughton, Shishir Prasad, Ganesh Krishnan, Rohit Deep, and Vijay Raghavendra. 2016. [Magellan: Toward building entity matching management systems over data science stacks](#). *Proc. VLDB Endow.*, 9(13):1581–1584.
- Bing Li, Yukai Miao, Yaoshu Wang, Yifang Sun, and Wei Wang. 2021a. Improving the efficiency and effectiveness for bert-based entity resolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13226–13233.
- Yuliang Li, Jinfeng Li, Yoshihiko Suhara, AnHai Doan, and Wang-Chiew Tan. 2020. [Deep entity matching with pre-trained language models](#). *Proceedings of the VLDB Endowment*, 14(1):50–60.
- Yuliang Li, Jinfeng Li, Yoshihiko Suhara, Jin Wang, Wataru Hirota, and Wang-Chiew Tan. 2021b. Deep entity matching: Challenges and opportunities. *Journal of Data and Information Quality (JDIQ)*, 13(1):1–17.
- Lassi A. Liikkanen and Antti Salovaara. 2015. [Music on youtube: User engagement with traditional, user-appropriated and derivative videos](#). *Computers in Human Behavior*, 50:108–124.
- Feng Liu, Deyi Tuo, Yinan Xu, and Xintong Han. 2023. [Coverhunter: Cover song identification with refined attention and alignments](#). *Preprint*, arXiv:2306.09025.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *Preprint*, arXiv:1907.11692.
- Claudio Lucchese, Franco Maria Nardini, Salvatore Orlando, Raffaele Perego, and Alberto Veneri. 2022. Ilmart: Interpretable ranking with constrained lambdamart. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2255–2259.
- Alberto Montero and Jorge Mora-Fernandez. 2020. Digital culture in youtube categories and interfaces: User experience and social interactions of the most popular videos and channels. In *International Conference on Human-Computer Interaction*, pages 383–401. Springer.
- Koji Okabe, Takafumi Koshinaka, and Koichi Shinoda. 2018. [Attentive statistics pooling for deep speaker embedding](#). In *Interspeech 2018*, interspeech₂₀₁₈.*ISCA*.
- OpenAI Inc. [ChatGPT](#).
- Matteo Paganelli, Donato Tiano, and Francesco Guerra. 2023. A multi-facet analysis of bert-based entity matching models. *The VLDB Journal*, pages 1–26.
- Ralph Peeters and Christian Bizer. 2022. [Supervised contrastive learning for product matching](#). In *Companion Proceedings of the Web Conference 2022*, WWW '22, page 248–251, New York, NY, USA. Association for Computing Machinery.
- Ralph Peeters and Christian Bizer. 2024. [Entity matching using large language models](#). *Preprint*, arXiv:2310.11244.
- Ralph Peeters, Reng Chiz Der, and Christian Bizer. 2023. Wdc products: A multi-dimensional entity matching benchmark. *arXiv preprint arXiv:2301.09521*.

- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2020. [Making monolingual sentence embeddings multilingual using knowledge distillation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4512–4525, Online. Association for Computational Linguistics.
- Jordan B. L. Smith, Masahiro Hamasaki, and Masataka Goto. 2017. [Classifying derivative works with search, text, audio and video features](#). In *2017 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1422–1427.
- Qiang Wu, Christopher J. C. Burges, Krysta M. Svore, and Jianfeng Gao. 2010. [Adapting boosting for information retrieval measures](#). 13(3):254–270.
- 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.
- Hong Xuan, Abby Stylianou, Xiaotong Liu, and Robert Pless. 2020. Hard negative examples are hard, but useful. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, pages 126–142. Springer.
- Dezhong Yao, Yuhong Gu, Gao Cong, Hai Jin, and Xinqiao Lv. 2022. [Entity resolution with hierarchical graph attention networks](#). In *Proceedings of the 2022 International Conference on Management of Data, SIGMOD '22*, page 429–442, New York, NY, USA. Association for Computing Machinery.
- Furkan Yesiler, Chris Tralie, Albin Correya, Diego F. Silva, Philip Tovstogan, Emilia Gómez, and Xavier Serra. 2019. Da-TACOS: A dataset for cover song identification and understanding. In *Proc. of the 20th Int. Soc. for Music Information Retrieval Conf. (ISMIR)*, pages 327–334, Delft, The Netherlands.
- Zhesong Yu, Xiaoshuo Xu, Xiaou Chen, and Deshun Yang. 2020. [Learning a representation for cover song identification using convolutional neural network](#). In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 541–545.