



Cover detection using dominant melody embeddings

Guillaume Doras, Geoffroy Peeters

► To cite this version:

Guillaume Doras, Geoffroy Peeters. Cover detection using dominant melody embeddings. ISMIR 2019, Nov 2019, Delft, Netherlands. hal-02457735

HAL Id: hal-02457735

<https://hal.science/hal-02457735>

Submitted on 28 Jan 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

COVER DETECTION USING DOMINANT MELODY EMBEDDINGS

Guillaume Doras

Sacem & Ircam Lab, CNRS, Sorbonne Université
guillaume.doras@sacem.fr

Geoffroy Peeters

LTCI, Telecom Paris, Institut Polytechnique de Paris
geoffroy.peeters@telecom-paris.fr

ABSTRACT

Automatic cover detection – the task of finding in an audio database all the covers of one or several query tracks – has long been seen as a challenging theoretical problem in the MIR community and as an acute practical problem for authors and composers societies. Original algorithms proposed for this task have proven their accuracy on small datasets, but are unable to scale up to modern real-life audio corpora. On the other hand, faster approaches designed to process thousands of pairwise comparisons resulted in lower accuracy, making them unsuitable for practical use.

In this work, we propose a neural network architecture that is trained to represent each track as a single embedding vector. The computation burden is therefore left to the embedding extraction – that can be conducted offline and stored, while the pairwise comparison task reduces to a simple Euclidean distance computation. We further propose to extract each track’s embedding out of its dominant melody representation, obtained by another neural network trained for this task. We then show that this architecture improves state-of-the-art accuracy both on small and large datasets, and is able to scale to query databases of thousands of tracks in a few seconds.

1. INTRODUCTION

Covers are different interpretations of the same original musical work. They usually share a similar melodic line, but typically differ greatly in one or several other dimensions, such as their structure, tempo, key, instrumentation, genre, etc. Automatic cover detection – the task of finding in an audio database all the covers of one or several query tracks – has long been seen as a challenging theoretical problem in MIR. It is also now an acute practical problem for copyright owners facing continuous expansion of user-generated online content.

Cover detection is not *stricto sensu* a classification problem: due to the ever growing amount of musical works (the classes) and the relatively small number of covers per work, the actual question is not so much “to which work this track belongs to ?” as “to which other tracks this track is the most similar ?”.

Formally, cover detection therefore requires to establish a similarity relationship S_{ij} between a query track A_i and a reference track B_j . It implies the composite of a feature extraction function f followed by a pairwise comparison function g , expressed as $S_{ij} = g(f(A_i), f(B_j))$. If f and g are independent, the feature extraction of the reference tracks B_j can be done offline and stored. The online feature extraction cost is then linear in the number of queries, while pairwise comparisons cost without optimisation scales quadratically in the number of tracks [16].

Efficient cover detection algorithms thus require a fast pairwise comparison function g . Comparing pairs of entire sequences, as DTW does, scales quadratically in the length of the sequences and becomes quickly prohibitive. At the opposite, reducing g to a simple Euclidean distance computation between tracks embeddings is independent of the length of the sequences. In this case, the accuracy of the detection entirely relies on the ability of f to extract the common musical facets between different covers.

In this work, we describe a neural network architecture mapping each track to a single embedding vector, and trained to minimize cover pairs Euclidean distance in the embeddings space, while maximizing it for non-cover pairs. We leverage on recent breakthroughs in dominant melody extraction, and show that the use of dominant melody embeddings yield promising performances both in term of accuracy and scalability.

The rest of the paper is organized as follow: we review in §2 the main concepts used in this work. We detail our method in §3, and describe and discuss in §4 and §5 the different experiments conducted and their results. We finally present a comparison with existing methods in §6. We conclude with future improvements to bring to our method.

2. RELATED WORK

We review here the main concepts used in this study.

2.1 Cover detection

Successful approaches in cover detection used an input representation preserving common musical facets between different versions, in particular dominant melody [19, 27, 40], tonal progression – typically a sequence of chromas [10, 12, 33, 39] or chords [2], or a fusion of both [11, 29]. Most of these approaches then computed a similarity score between pairs of melodic and/or harmonic sequences, typically a cross-correlation [10], a variant of the DTW algorithm [12, 20, 33, 39], or a combination of both [25].



© Guillaume Doras, Geoffroy Peeters. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0).
Attribution: Guillaume Doras, Geoffroy Peeters. “Cover Detection Using Dominant Melody Embeddings”, 20th International Society for Music Information Retrieval Conference, Delft, The Netherlands, 2019.

These approaches lead to good results when evaluated on small datasets – at most a few hundreds of tracks, but are not scalable beyond due to their expensive comparison function. Faster methods have recently been proposed, based on efficient comparison of all possible subsequences pairs between chroma representations [34], or similarity search between 2D-DFT sequences derived from CQTs overlapping windows [31], but remain too costly to be scalable to query large modern audio databases.

Another type of method has been proposed to alleviate the cost of the comparison function and to shift the burden to the audio features extraction function – which can be done offline and stored. The general principle is to encode each audio track as a single scalar or vector – its embedding – and to reduce the similarity computation to a simple Euclidean distance between embeddings. Originally, embeddings were for instance computed as a single hash encoding a succession of pitch landmarks [3], or as a vector obtained by PCA dimensionality reduction of a chromagram’s 2D-DFT [4] or with locality-sensitive hashing of melodic excerpts [19].

As for many other MIR applications, ad-hoc – and somewhat arbitrary – hand-crafted features extraction was progressively replaced with data-driven automatic feature learning [15]. Different attempts to learn common features between covers have since been proposed: in particular, training a k -means algorithm to learn to extract an embedding out of chromagram’s 2D-DFT lead to significant results improvements on large datasets [16]. Similar approaches, commonly referred to as *metric learning* approaches, have been used in different MIR contexts, such as music recommendation [21, 41], live song identification [38], music similarity search [24], and recently cover detection [23].

2.2 Metric learning

Although the concept can be traced back to earlier works [1, 8], the term of metric learning was probably coined first in [43] to address this type of clustering tasks where the objective is merely to assess whether different samples are similar or dissimilar. It has since been extensively used in the image recognition field in particular [14, 36, 37].

The principle is to learn a mapping between the input space and a latent manifold where a simple distance measure (such as Euclidean distance) should approximate the neighborhood relationships in the input space. There is however a trivial solution to the problem, where the function ends up mapping all the examples to the same point. Contrastive Loss was introduced to circumvent this problem, aiming at simultaneously *pulling* similar pairs together and *pushing* dissimilar pairs apart [13].

However, when the amount of labels becomes larger, the number of dissimilar pairs becomes quickly intractable. It was moreover observed in practice that once the network has become reasonably good, negative pairs become relatively easy to discern, which stalls the training of the discriminative model. *Pair mining* is the strategy of training the model only with hard pairs, i.e. positive (resp. nega-

tive) pairs with large (resp. small) distances [35]. Further improvement was introduced with the triplet loss, which is used to train a model to map each sample to an embedding that is closer to all of its positive counterparts than it is to all of its negative counterparts [30]. Formally, for all triplets $\{a, p, n\}$ where a is an anchor, and p or n is one of its positive or negative example, respectively, the loss to minimize is expressed as $\ell = \max(0, d_{ap} + \alpha - d_{an})$, where α is a margin and d_{ap} and d_{an} are the distances between each anchor a and p or n , respectively.

2.3 Dominant melody extraction

Dominant melody extraction has long been another challenging problem in the MIR community [18, 28, 42]. A major breakthrough was brought recently with the introduction of a convolutional network that learns to extract the dominant melody out of the audio Harmonic CQT [7]. The HCQT is an elegant and astute representation of the audio signal in 3 dimensions (time, frequency, harmonic), stacking along the third dimension several standard CQTs computed at different minimal multiple frequencies. Harmonic components of audio signal will thus be represented along the third dimension and be localized at the same location along the first and second dimensions. This representation is particularly suitable for melody detection, as it can be directly processed by convolutional networks, whose 3-D filters can be trained to localize in the time and frequency plan the harmonic components.

In a recent work [9], we suggested in an analogy with image processing that dominant melody extraction can be seen as a type of image segmentation, where contours of the melody have to be isolated from the surrounding background. We have thus proposed for dominant melody estimation an adaptation of U-Net [26] – a model originally designed for medical image segmentation – which slightly improves over [7].

3. PROPOSED METHOD

We present here the input data used to train our network, the network architecture itself and its training loss.

3.1 Input data

We have used as input data the dominant melody 2D representation (F0-CQT) obtained by the network we proposed in [9]. The frequency and time resolutions required for melody extraction (60 bins per octave and 11 ms per time frame) are not needed for cover detection. Moreover, efficient triplet loss training requires large training batches, as we will see later, so we reduced data dimensionality as depicted on Figure 2.

The F0-CQT is **a**) trimmed to keep only 3 octaves around its mean pitch (180 bins along the frequency axis), and only the first 3 minutes of the track (15500 time frames) – if shorter, the duration is not changed. The resulting matrix is then **b**) downsampled via bilinear 2D interpolation with a factor 5. On the frequency axis, the semi-tone resolution is thus reduced from five to one bin,




Figure 1: Convolutional model (time on the first dimension, frequency on the second dimension).

which we considered adequate for cover detection. On the time axis, it is equivalent to a regular downsampling.

Finally, as the representation of different tracks with possibly different durations shall be batched together during training, the downsampled F0-CQT is c) shrunk or stretched along the time axis by another bilinear interpolation to a fixed amount of bins (1024). This operation is equivalent to a tempo change: for the 3 minutes trimmed, shrinking is equivalent to multiply the tempo by a factor 3. We argue here that accelerated or decelerated version of a cover is still a cover of the original track.




Figure 2: Input data pre-processing: dominant melody is extracted from HCQT, then a) F0 output is trimmed, b) downsampled by a factor 5 and c) resized time-wise to 1024 bins (time bins on first dimension, frequency bins on second dimension).

3.2 Model

The proposed model is a simple convolutional network pictured in Figure 1. As we are constrained by the input data shape, whose time dimension is much larger than its frequency dimension, only five layers blocks are needed. Each layer block consists of a batch normalization layer, a convolution layer with 3×3 kernels and a mean-pooling layer with a 3×2 kernel and 3×2 stride in order to reduce time dimensionality faster than frequency dimensionality. A dropout rate of 0.1, 0.1, 0.2 and 0.3 is applied to the blocks 2, 3, 4 and 5, respectively.

The first convolutional layer has K kernels, and this number is doubled at each level (i.e. the deeper layer outputs $2^4 K$ -depth tensors). The penultimate layer averages along frequency and time axes to obtain a vector. A last dense layer outputs and L2-normalizes the final embedding vector of size E .

Our assumption behind the choice of this convolutional architecture is that we expect it to learn similar patterns in the dominant melody, at different scales (tempo invariance) and locations (key and structure invariance).

3.3 Objective loss

We use a triplet loss with online semi-hard negative pairs mining as in [30]. In practice, triplet mining is done within each training batch: instead of using all possible triplets, each track in the batch is successively considered as the anchor, and compared with all its covers in the batch. For each of these positives pairs, if there are negatives such as $d_{an} < d_{ap}$, then only the one with the highest d_{an} is kept. If no such negative exist, then only the one with the lowest d_{an} is kept. Other negatives are not considered.

Model is fit with Adam optimizer [17], with initial learning rate at $1e^{-4}$, divided by 2 each time the loss on the evaluation set does not decrease after 5k training steps. Training is stopped after 100k steps, or if the learning rate falls below $1e^{-7}$. The triplet loss was computed using squared Euclidean distances (i.e. distances are within the $[0, 4]$ range), and the margin was set to $\alpha = 1$.

3.4 Dataset

As metric learning typically requires large amount of data, we fetched from internet the audio of cover tracks provided by the SecondHandSongs website API¹. Only works with 5 to 15 covers, and only tracks lasting between 60 and 300 seconds were considered, for a total of $W = 7460$ works and $T = 62310$ tracks.

The HCQT was computed for those 62310 tracks as detailed in [7], i.e. with $f_{min} = 32.7$ Hz and 6 harmonics. Each CQT spans 6 octaves with a resolution of 5 bins per semi-tone, and a frame duration of ~ 11 ms. The implementation was done with the Librosa library [22].

The dominant melody was extracted for these 62310 HCQT with the network we described in [9], and the output was trimmed, downsampled and resized as described in §3.1.

4. PRELIMINARY EXPERIMENTS

We present here some experiments conducted to develop the system. The 7460 works were split into disjoint train and evaluation sets, with respectively 6216 and 1244 works and five covers per work. The evaluation set represents $\sim 20\%$ of the training set, which we considered fair enough given the total amount of covers. The same split has been used for all preliminary experiments.

¹ <https://secondhandsongs.com/>

4.1 Metrics

Ideally, we expect the model to produce embeddings such that cover pair distances are low and non-cover pair distances are high, with a large gap between the two distributions. In the preliminary experiments, we have thus evaluated the separation of the cover pairs distance distribution $p_c(d)$ from the non-cover pairs distance distribution $p_{nc}(d)$ with two metrics:

- the ROC curve plots the true positives rate (covers, TPR) versus the false positive rate (non-covers, FPR) for different distance d thresholds. We report the area under the ROC curve (AuC), which gives a good indication about the distributions separation. We also report the TPR corresponding to an FPR of 5% (TPR@5%), as it gives an operational indication about the model's discriminative power.
- we also report the Bhattacharyya coefficient (BC), expressed as $\sum_d \sqrt{p_c(d)p_{nc}(d)}$, as it directly measures the separation between the distributions (smaller is better) [6].

4.2 Influence of input data

We first compared the results obtained for different inputs data: chromas and CQT computed using Librosa [22], and the dominant melody computed as described in 3.1. As shown on Figure 3 (left), dominant melody yields the best results. It does not imply that melody features are more suited than tonal features for cover detection, but shows that convolutional kernels are better at learning similar patterns at different scales and locations across different tracks when the input data is sparse, which is not the case for chromas and CQT.




Figure 3: Scores obtained on evaluation set for a model trained - with chromas, CQT or F0 (left). - on the F0 with various octaves spans (middle). - on the F0 with various durations (right).

Results obtained when trimming the F0-CQT with various octaves and time spans are also shown Figure 3. It appears that keeping 3 octaves around the mean pitch of the dominant melody and a duration of 2 to 3 minutes yields the best results. Smaller spans do not include enough information, while larger spans generate confusion.

All other results presented below are thus obtained with the dominant melody 2D representation as input data, and a span of 3 octaves and 180 seconds for each track.

4.3 Influence of model and training parameters

We then compared the results obtained for different numbers of kernels in the first layer (K) and the corresponding sizes of the embeddings (E). As shown on Figure 4

(left), results improve for greater K, which was expected. However, increasing K above a certain point does not improve the results further, as the model has probably already enough freedom to encode common musical facets.




Figure 4: Scores obtained on evaluation set for a model trained - with various K/E (left) - with various B (right).

We have then compared the results obtained for different sizes of training batches (B). As shown on Figure 4 (right), results improve with larger B: within larger batches, each track will be compared with a greater number of non-covers, improving the separation between clusters of works. A closer look at the distances shows indeed that the negative pairs distance distribution $p_{nc}(d)$ gets narrower for larger batches (not showed here). Due to GPU memory constraints, we have not investigated values above $B=100$.

All other results presented below are obtained with $K=64$, $E=512$ and $B=100$.

5. LARGE SCALE LOOKUP EXPERIMENTS

We now present experiments investigating the realistic use case, i.e. large audio collections lookup. When querying an audio collection, each query track can be of three kinds: **a)** it is already present in the database, **b)** it is a cover of some other track(s) already in the database, or **c)** it is a track that has no cover in the database. The case **a)** corresponds to the trivial case, where the query will produce a distance equal to zero when compared with itself, while case **c)** corresponds to the hard case where neither the query or any cover of the query have been seen during training. We investigate here the case **b)**, where the query track itself has never been seen during training, but of which at least one cover has been seen during training.

5.1 Metrics

In these experiments, we are interested in measuring our method's ability to find covers in the reference set when queried with various unknown tracks. This is commonly addressed with the metrics proposed by MIREX² for the cover song identification task: the mean rank of first correct result (MR1), the mean number of true positives in the top ten positions (MT10) and the Mean Average Precision (MAP). We refer the reader to [32] for a detailed review of these standard metrics. We also report here the TPR@5%, already used in the preliminary experiments.

² <https://www.music-ir.org/mirex/wiki/2019>

5.2 Structuring the embeddings space

We study here the role of the training set in structuring the embeddings space, and in particular the role of the number of covers of each work. More precisely, we tried to show evidence of the *pushing* effect (when a query is pushed away from all its non-covers clusters) and the *pulling* effect (when a query is pulled towards its unique covers cluster).

To this aim, we built out of our dataset a query and a reference set. The query set includes 1244 works with five covers each. The reference set includes P of the remaining covers for each of the 1244 query works, and N covers for each other work not included in the query set (Figure 5).




Figure 5: Big dotted oval represents the embeddings space \mathcal{E} . Smaller dotted ovals represent work clusters of covers. Red crossed circles represent the positions in the manifold of the reference tracks w_j that are not covers of query track q_i (N per cluster). Green circles represent the positions of the tracks that are covers of query track q_i (P per cluster).

Pushing covers We first train our model on the reference set with fixed $P=5$. We compute query tracks embeddings with the trained model, compute pairwise distances between query and reference embeddings, as well as the different metrics. We repeat this operation for different values of $N \in [2, \dots, 10]$, and report results on Figure 6 (left). We report MR1's percentile (defined here as MR1 divided by the total of reference tracks, in percent) instead of MR1, because the number of reference tracks varies with N .




Figure 6: Query scores for different training/reference sets. Left: increasing N , number of covers of non-queries, P fixed. Right: increasing P , number of covers of queries, N fixed. (Y-axes scale MAP and TPR on the left, and the MR1 percentile on the right).

The MAP only slightly decreases as N increases, which indicates that the precision remains stable, even though the number of examples to sort and to rank is increasing. Moreover, the MR1 percentile and the TPR@5% clearly improve as N increases. As P is fixed, it means that the ranking and the separation between covers and non-covers clusters is improving as the non-queries clusters are consolidated, which illustrates the expected *pushing* effect.

Pulling covers We reproduce the same protocol again, but now with $N=5$ fixed and for different values of $P \in [2, \dots, 10]$. We report results on Figure 6 (right). It appears clearly that all metrics improve steadily as P increases, even though the actual query itself has never been seen during training. As N is fixed, this confirms the intuition that the model will get better in locating unseen tracks closer to their work's cluster if trained with higher number of covers of this work, which illustrates the expected *pulling* effect.

5.3 Operational meaning of $p_c(d)$ and $p_{nc}(d)$

We now investigate further the distance distributions of cover and non-cover pairs. To this aim, we randomly split our entire dataset into a query and a reference set with a 1:5 ratio (resp. 10385 and 51925 tracks). Query tracks are thus not seen during training, but might have zero or more covers in the reference set.

Covers probability Computing queries vs. references pairwise distances gives the distributions $p_c(d)$ and $p_{nc}(d)$ shown on Figure 7 (left). Using Bayes' theorem, it is straightforward to derive from $p_c(d)$ and $p_{nc}(d)$ the probability for a pair of tracks to be covers given their distance d (Figure 7, right). This curve has an operational meaning, as it maps a pair's distance with a probability of being covers without having to rank it among the entire dataset.

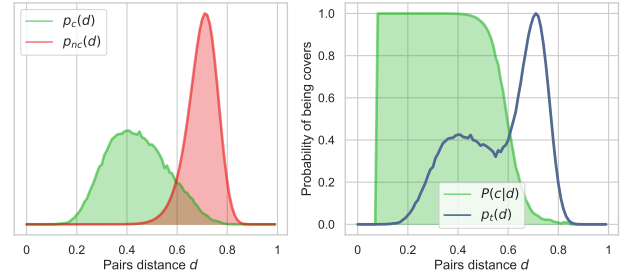


Figure 7: Left: separation of $p_c(d)$ (green) and $p_{nc}(d)$ (red). Right: probability of being a covers pair given the distance d (green) and total pairs distance distribution $p_c(d) + p_{nc}(d)$ (blue).

Easy and hard covers We repeat the previous test five times with random splits, and report metrics in Table 1. At first sight, MR1 and MT@10 could seem inconsistent, but a closer look at the results gives an explanation. To illustrate what happens, imagine a set of five queries where the first query ranks ten covers correctly in the first ten positions, e.g. because they are all very similar, while all other four queries have their first correct answer at rank 100. This would yield to $MT@10=2.0$, and $MR1=80.2$. This kind of discrepancy between MR1 and MT@10 reflects the fact that some works in our dataset have similar covers that are easily clustered, while other are much more difficult to discriminate. This can be observed on the positive pairs distribution $p_c(d)$ on Figure 7 (left), which is spread over a large range of distances.

	MAP	MT@10	MR1	BC	AuC	TPR@5%
Proposed	0.39 (<0.01)	2.90 (0.03)	581 (59)	0.40 (<0.01)	0.97 (<0.01)	0.88 (<0.01)

Table 1: Results of query set lookup in reference set.

6. COMPARISON WITH OTHER METHODS

6.1 Comparison on small dataset

We first compared with two recent methods [31, 34], who reported results for a small dataset of 50 works with 7 covers each. The query set includes five covers of each work (250 tracks), while the reference set includes each work’s remaining two covers (100 tracks). As this dataset is not publicly available anymore, we have mimicked it extracting randomly 350 tracks out of own dataset³.

Our data-driven model can however not be trained with only 100 tracks of the reference set, as it would overfit immediately. We have thus trained our model on our full dataset, with two different setups: **a)** excluding the 350 tracks reserved for the query and reference sets. **b)** excluding the 250 tracks of the query set, but including the 100 tracks of the reference set. We repeated this operation ten times for each setup, and report the mean and standard deviation on Table 2 for the same metrics used in [31, 34], as well as the p-value obtained by a statistical significance t-test carried out on results series.

	MAP	P@10	MR1
[34]	0.591	0.140	7.910
[31]	0.648	0.145	8.270
Proposed a)	0.675	0.165	3.439
	(0.04), p=.29	(0.005), p<.001	(1.062), p<.001
Proposed b)	0.782	0.179	2.618
	(0.104), p<.01	(0.014), p<.001	(1.351), p<.001

Table 2: Comparison between recent method [31, 34] and our proposed method on a small dataset (precision at 10 P@10 is reported instead of MT@10. As there are only two covers per work in the reference set, P@10 maximum value is 0.2).

Our method significantly improve previous results: for the hardest case **a)** where the model has not seen any queries work during training, embeddings space has been sufficiently structured to discriminate the unseen works from the other training clusters (*pushing* effect). For the easier case **b)**, the *pulling* effect from the known queries covers provides further improvement.

6.2 Comparison on large dataset

We also compared with [16], who is to our knowledge the last attempt to report results for thousands of queries and references – a more realistic use case. This paper reported results on the SecondHandSong (SHS) subset of the MillionSong dataset (MSD) [5] for two experiments: **a)** only the training set of 12960 covers of 4128 works was used both as the query and reference sets. **b)** the SHS MSD test set of 5236 covers of 1726 works was used to query the entire MSD used as reference.

The SHS MSD is not available anymore. However, as our dataset has also been built from the SHS covers list, we consider that results can be compared³. We have therefore randomly generated out of our dataset a training and a test

set mimicking the original ones. We trained our model on the training set, and perform the pairwise distances computation between the query and reference sets (as the query set is included in the reference set, we excluded for comparison the pairs of the same track). For experiment **b)**, we have used our entire dataset as reference set as we do not have one million songs. We have repeated this operation five times and report in Table 3 the mean and standard deviations for the same metrics used in [16], as well as MR1, MT@10 and the p-value of the t-test carried out.

		MAP	MR	MT@10	MR1
a)	[16]	0.285	1844	-	-
	Proposed	0.936	78	2.010	33
		(0.001), p<.001	(6), p<.001	(<0.001)	(3)
b)	[16]	0.134	173117	-	-
	Proposed	0.220	3865	1.622	430
		(0.007), p<.001	(81), p<.001	(0.003)	(19)

Table 3: Comparison between method [16] and our proposed method on a large dataset (MR=Mean rank). For **b)**, the MR percentile should be compared, as our reference set does not have 1M tracks (6th vs. 17th percentile for [16]).

Our method significantly improve previous results. For case **a)**, results are notably good, which is not surprising as the model has already seen all the queries during the training. Case **b)** is on the other hand the hardest possible configuration, where the model has not seen any covers of the queries works during training, and clusterisation of unseen tracks entirely relies on the *pushing* effect.

As to our method’s computation times, we observed on a single Nvidia GPU Titan XP for a ~3 mn audio track: ~10 sec for F0 extraction, ~1 sec for embeddings computation, and less than 0.2 sec for distances computation with the full dataset embeddings (previously computed offline).

7. CONCLUSION

In this work, we presented a method for cover detection, using a convolutional network which encodes each track as a single vector, and is trained to minimize cover pairs Euclidean distance in the embeddings space, while maximizing it for non-covers. We show that extracting embeddings out of the dominant melody 2D representation drastically yields better results compared to other spectral representations: the convolutional model learns to identify similar patterns in the dominant melody at different scales and locations (tempo, key and structure invariance).

We have also shown that our method scales to audio databases of thousands of tracks. Once trained for a given database, it can be used to assess the probability for an unseen track to be a cover of any known track without having to be compared to the entire database. We have finally shown that our method improves previous methods both on small and large datasets.

In the future, we plan to grow our training dataset to address the realistic use case where collections of millions of tracks should be queried: as for many other data-driven problems, will the cover detection problem be solved if the embeddings space is sufficiently structured?

³ We acknowledge that, strictly speaking, results can not be directly compared to rank methods, as original datasets of [34], [31] and [16] are not available any longer. They however reflect the level of performance of the proposed method.

Acknowledgements.

The dominant melody representations of tracks used as training dataset in this work is available upon request.

We are grateful to Xavier Costaz at Sacem for his suggestions and support, and to Philippe Esling at Ircam for fruitful discussions.

8. REFERENCES

- [1] Pierre Baldi and Yves Chauvin. Neural networks for fingerprint recognition. *Neural Computation*, 5(3):402–418, 1993.
- [2] Juan Pablo Bello. Audio-based cover song retrieval using approximate chord sequences: Testing shifts, gaps, swaps and beats. In *Proceedings of ISMIR (International Society of Music Information Retrieval)*, 2007.
- [3] Thierry Bertin-Mahieux and Daniel PW Ellis. Large-scale cover song recognition using hashed chroma landmarks. In *Proceedings of IEEE WASPAA (Workshop on Applications of Signal Processing to Audio and Acoustics)*, pages 117–120. IEEE, 2011.
- [4] Thierry Bertin-Mahieux and Daniel PW Ellis. Large-scale cover song recognition using the 2d fourier transform magnitude. In *Proceedings of ISMIR (International Society of Music Information Retrieval)*, 2012.
- [5] Thierry Bertin-Mahieux, Daniel PW Ellis, Brian Whitman, and Paul Lamere. The million song dataset. *Proceedings of ISMIR (International Society of Music Information Retrieval)*, 2011.
- [6] Anil Bhattacharyya. On a measure of divergence between two statistical populations defined by their probability distributions. *Bull. Calcutta Math. Soc.*, 35:99–109, 1943.
- [7] Rachel M Bittner, Brian McFee, Justin Salamon, Peter Li, and Juan P Bello. Deep salience representations for f0 estimation in polyphonic music. In *Proceedings of ISMIR (International Society of Music Information Retrieval)*, 2017.
- [8] Jane Bromley, Isabelle Guyon, Yann LeCun, Edward Säckinger, and Roopak Shah. Signature verification using a "siamese" time delay neural network. In *Advances in Neural Information Processing Systems*, pages 737–744, 1994.
- [9] Guillaume Doras, Philippe Esling, and Geoffroy Peeters. On the use of u-net for dominant melody estimation in polyphonic music. In *International Workshop on Multilayer Music Representation and Processing (MMRP)*, pages 66–70. IEEE, 2019.
- [10] Daniel PW Ellis and Graham E Poliner. Identifying-cover songs with chroma features and dynamic programming beat tracking. In *Proceedings of ICASSP (International Conference on Acoustics, Speech and Signal Processing)*. IEEE, 2007.
- [11] Rémi Foucard, Jean-Louis Durrieu, Mathieu Lagrange, and Gaël Richard. Multimodal similarity between musical streams for cover version detection. In *Proceedings of ICASSP (International Conference on Acoustics, Speech and Signal Processing)*. IEEE, 2010.
- [12] Emilia Gómez and Perfecto Herrera. The song remains the same: identifying versions of the same piece using tonal descriptors. In *Proceedings of ISMIR (International Society of Music Information Retrieval)*, 2006.
- [13] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *Proceedings of IEEE CVPR (Conference on Computer Vision and Pattern Recognition)*, volume 2, pages 1735–1742. IEEE, 2006.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016.
- [15] Eric J Humphrey, Juan Pablo Bello, and Yann LeCun. Moving beyond feature design: Deep architectures and automatic feature learning in music informatics. In *Proceedings of ISMIR (International Society of Music Information Retrieval)*, 2012.
- [16] Eric J Humphrey, Oriol Nieto, and Juan Pablo Bello. Data driven and discriminative projections for large-scale cover song identification. In *Proceedings of ISMIR (International Society of Music Information Retrieval)*, 2013.
- [17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [18] Anssi Klapuri. Multiple fundamental frequency estimation by summing harmonic amplitudes. In *Proceedings of ISMIR (International Society of Music Information Retrieval)*, 2006.
- [19] Matija Marolt. A mid-level representation for melody-based retrieval in audio collections. *IEEE Transactions on Multimedia*, 10(8):1617–1625, 2008.
- [20] Benjamin Martin, Daniel G Brown, Pierre Hanna, and Pascal Ferraro. Blast for audio sequences alignment: a fast scalable cover identification. In *Proceedings of ISMIR (International Society of Music Information Retrieval)*, 2012.
- [21] Brian McFee, Luke Barrington, and Gert Lanckriet. Learning content similarity for music recommendation. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(8):2207–2218, 2012.
- [22] Brian McFee, Colin Raffel, Dawen Liang, Daniel PW Ellis, Matt McVicar, Eric Battenberg, and Oriol Nieto. librosa: Audio and music signal analysis in python. In *Proceedings of the 14th python in science conference*, pages 18–25, 2015.

- [23] Xiaoyu Qi, Deshun Yang, and Xiaoou Chen. Triplet convolutional network for music version identification. In *International Conference on Multimedia Modeling*, pages 544–555. Springer, 2018.
- [24] Colin Raffel and Daniel PW Ellis. Pruning subsequence search with attention-based embedding. In *Proceedings of ICASSP (International Conference on Acoustics, Speech and Signal Processing)*. IEEE, 2016.
- [25] Suman Ravuri and Daniel PW Ellis. Cover song detection: from high scores to general classification. In *Proceedings of ICASSP (International Conference on Acoustics, Speech and Signal Processing)*. IEEE, 2010.
- [26] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [27] Christian Sailer and Karin Dressler. Finding cover songs by melodic similarity. *MIREX extended abstract*, 2006.
- [28] Justin Salamon and Emilia Gómez. Melody extraction from polyphonic music signals using pitch contour characteristics. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(6):1759–1770, 2012.
- [29] Justin Salamon, Joan Serra, and Emilia Gómez. Tonal representations for music retrieval: from version identification to query-by-humming. *International Journal of Multimedia Information Retrieval*, 2(1):45–58, 2013.
- [30] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of IEEE CVPR (Conference on Computer Vision and Pattern Recognition)*, pages 815–823, 2015.
- [31] Prem Seetharaman and Zafar Rafii. Cover song identification with 2d fourier transform sequences. In *Proceedings of ICASSP (International Conference on Acoustics, Speech and Signal Processing)*. IEEE, 2017.
- [32] Joan Serrà. *Music similarity based on sequences of descriptors tonal features applied to audio cover song identification*. PhD thesis, Universitat Pompeu Fabra, Spain, 2007.
- [33] Xavier Serra, Ralph G Andrzejak, et al. Cross recurrence quantification for cover song identification. *New Journal of Physics*, 11(9):093017, 2009.
- [34] Diego F Silva, Chin-Chin M Yeh, Gustavo Enrique de Almeida Prado Alves Batista, Eamonn Keogh, et al. Simple: assessing music similarity using subsequences joins. In *Proceedings of ISMIR (International Society of Music Information Retrieval)*, 2016.
- [35] Edgar Simo-Serra, Eduard Trulls, Luis Ferraz, Iasonas Kokkinos, Pascal Fua, and Francesc Moreno-Noguer. Discriminative learning of deep convolutional feature point descriptors. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 118–126, 2015.
- [36] Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *Advances in Neural Information Processing Systems*, pages 1857–1865, 2016.
- [37] Hyun Oh Song, Yu Xiang, Stefanie Jegelka, and Silvio Savarese. Deep metric learning via lifted structured feature embedding. In *Proceedings of IEEE CVPR (Conference on Computer Vision and Pattern Recognition)*, pages 4004–4012. IEEE, 2016.
- [38] TJ Tsai, Thomas Prätzlich, and Meinard Müller. Known artist live song id: A hashprint approach. In *Proceedings of ISMIR (International Society of Music Information Retrieval)*, 2016.
- [39] Wei-Ho Tsai, Hung-Ming Yu, Hsin-Min Wang, et al. Query-by-example technique for retrieving cover versions of popular songs with similar melodies. In *Proceedings of ISMIR (International Society of Music Information Retrieval)*, 2005.
- [40] Wei-Ho Tsai, Hung-Ming Yu, Hsin-Min Wang, and Jorng-Tzong Horng. Using the similarity of main melodies to identify cover versions of popular songs for music document retrieval. *Journal of Information Science & Engineering*, 24(6), 2008.
- [41] Aaron Van den Oord, Sander Dieleman, and Benjamin Schrauwen. Deep content-based music recommendation. In *Advances in neural information processing systems*, pages 2643–2651, 2013.
- [42] Emmanuel Vincent, Nancy Bertin, and Roland Badeau. Adaptive harmonic spectral decomposition for multiple pitch estimation. *IEEE Transactions on Audio, Speech and Language Processing*, 18(3):528–537, 2010.
- [43] Eric P Xing, Michael I Jordan, Stuart J Russell, and Andrew Y Ng. Distance metric learning with application to clustering with side-information. In *Advances in Neural Information Processing Systems*, pages 521–528, 2003.