

# Cover Song Identification with Timbral Shape Sequences

Christopher J. Tralie  
chris.tralie@gmail.com

Paul Bendich  
bendich@math.duke.edu

July 17, 2015

## Abstract

We introduce a novel low level feature for identifying cover songs which quantifies the relative changes in the smoothed frequency spectrum of a song. Our key insight is that a sliding window representation of a chunk of audio can be viewed as a time-ordered point cloud in high dimensions. For corresponding chunks of audio between different versions of the same song, these point clouds are approximately rotated, translated, and scaled copies of each other. If we treat MFCC embeddings as point clouds and cast the problem as a relative shape sequence, we are able to correctly identify 42/80 cover songs in the “Covers 80” dataset. By contrast, all other work to date on cover songs exclusively relies on matching note sequences from Chroma derived features.

## 1 Introduction

Automatic cover song identification is a surprisingly difficult classical problem that has long been of interest to the music information retrieval community [5]. This problem is significantly more challenging than traditional audio fingerprinting because a combination of tempo changes, musical key transpositions, embellishments in time and expression, and changes in vocals and instrumentation can all occur simultaneously between the original version of a song and its cover. Hence, low level features used in this task need to be robust to all of these phenomena, ruling out raw forms of popular features such as MFCC, CQT, and Chroma.

One prior approach, as reviewed in Section 2, is to compare beat-synchronous sequences of chroma vectors between candidate covers. The beat-syncing helps this be invariant to tempo, but it is still not invariant to key. However, many schemes have been proposed to deal with this, up to and including a brute force check over all key transpositions.

Chroma representations factor out some timbral information by folding together all octaves, which is sensible given the effect that different instruments and recording environments have on timbre. However, valuable non-pitch information which is preserved between cover versions, such as spectral fingerprints from drum patterns, is obscured in Chroma representation. This motivated us to take another look at whether timbral-based features could be used at all for this problem. Our idea is that even if absolute timbral information is vastly different between two versions of the same song, the *relative evolution* of timbre over time should be comparable.

With careful centering and normalization within small windows to combat differences in global timbral drift between the two songs, we are indeed able to design shape features which are approximately invariant to cover. These features, which are based on self-similarity matrices of MFCC coefficients, can be used on their own to effectively score cover songs. This, in turn, demonstrates that even if absolute pitch is obscured and blurred, cover song identification is still possible.

Section 2 reviews prior work in cover song identification. Our method is described in detail by Sections 3 and 4. Finally, we report results on the “Covers 80” benchmark dataset [7] in Section 5, and we apply our algorithm to the recent “Blurred Lines” copyright controversy.

## 2 Prior Work

To the best of our knowledge, all prior low level feature design for cover song identification has focused on Chroma-based representations alone. The cover songs problem statement began with the work of [5],

which used FFT-based cross-correlation of all key transpositions of beat-synchronous chroma between two songs. A follow-up work [8] showed that high passing such cross-correlation can lead to better results. In general, however, cross-correlation is not robust to changes in timing, and it is also a global alignment technique. Serra [22] extended this initial work by considering dynamic programming local alignment of chroma sequences, with follow-up work and rigorous parameter testing and an “optimal key transposition index” estimation presented in [23]. The same authors also showed that a delay embedding of statistics spanning multiple beats before local alignment improves classification accuracy [25]. In a different approach, [14] compared modeled covariance statistics of all chroma bins, as well as comparing covariance statistics for all pairwise differences of beat-level chroma features, which is not unlike the “bag of words” and bigram representations, respectively, in text analysis. Other work tried to model sequences of chords [2] as a slightly higher level feature than chroma. Slightly later work concentrated on fusing the results of music separated into melody and accompaniment [11] and melody, bass line, and harmony [21], showing improvements over matching chroma on the raw audio. The most recent work on cover song identification has focused on fast techniques for large scale pitch-based cover song identification, using a sparse set of approximate nearest neighbors [28] and low dimensional projections [12]. Authors in [9] and [17] also use the *magnitude* of the 2D Fourier Transform of a sequences of chroma vectors treated as an image, so the resulting coefficients will be automatically invariant to key and time shifting without any extra computation, at the cost of some discriminative power.

Outside of cover song identification, there are other works which examine gappy sequences of MFCC in music, such as [4]. However, these works look at matched sequences of MFCC-like features in their original feature space. By contrast, in our work, we examine the *relative* shape of such features. Finally, we are not the first to consider shape in an applied musical context. For instance, [29] turns sequences of notes in sheet music into plane curves, whose curvature is then examined. To our knowledge, however, we are the first to explicitly model shape in musical audio for version identification.

### 3 Time Ordered Point Clouds from Blocks of Audio

The first step of our algorithm uses a timbre-based method to turn a block of audio into what we call a *time-ordered point cloud*. We can then compare to other time-ordered point clouds in a rotation, translation, and scale invariant manner using normalized Euclidean Self-Similarity matrices (Section 3.3). The goal is to then match up the relative shape of musical trajectories between cover versions.

#### 3.1 Point Clouds from Blocks and Windows

We start with a song, which is a function of time  $f(t)$  that has been discretized as some vector  $X$ . In the following discussion, the symbol  $X(a, b)$  means the song portion beginning at time  $t = a$  and ending at time  $t = b$ . Given  $X$ , there are many ways to summarize a chunk of audio  $w \in X$ , which we call a *window*, as a point in some feature space. We use the classical Mel-Frequency Cepstral coefficient representation [3], which is based on a perceptually motivated log frequency and log power short-time Fourier transform that preserves timbral information. In our application, we perform an MFCC with 20 coefficients, giving rise to a 20-dimensional point.

$$MFCC(w) \in \mathbb{R}^{20} \tag{1}$$

Given a longer chunk of audio, which we call a *block*, we can use the above embedding on a collection of  $K$  windows that cover the block to construct a collection of points, or a *point cloud*, representing that block. More formally, given a block covering a range  $[t_1, t_2]$ , we want a set of window intervals  $[a_i, b_i]$ , with  $i = 1..K$ , so that

- $a_i < b_i$
- $a_i < a_{i+1}, b_i < b_{i+1}$
- $\cup_{i=1}^K [a_i, b_i] = [t_1, t_2]$

Where  $t_1$ ,  $t_2$ ,  $a_i$ , and  $b_i$  are all discrete time indices into the sampled audio  $X$ . Hence, our final operator takes a set of time-ordered intervals  $\{[a_1, b_1], [a_2, b_2], \dots, [a_K, b_K]\}$  which cover a block  $[t_1, t_2]$  and turns them into a  $K$ -dimensional point cloud in  $\mathbb{R}^{20}$

$$PC(\{[a_1, b_1], \dots, [a_K, b_K]\}) = \{MFCC(X(a_1, b_1)), \dots, MFCC(X(a_K, b_K))\} \quad (2)$$

### 3.2 Beat-Synchronous Blocks

As many others in the MIR community have done, including [5] and [8] for the cover songs application, we compute our features synchronized within beat intervals. We use a simple dynamic programming beat tracker developed in [6]. Similarly to [8], we bias the beat tracker with three initial tempo levels: 60BPM, 120BPM, and 180BPM, and we compare the embeddings from all three levels against each other when comparing two songs, taking the best score out of the 9 combinations. This is to mitigate the tendency of the beat tracker to double or halve the true beat intervals of different versions of the same song when there are tempo changes between the two. The trade-off is of course additional computation. We should note that other cover song works, such as [23], avoid beat tracking step altogether, hence bypassing these problems. However, it is important for us to align our sequences as well as possible in time so that shape features are in correspondence, and this is a straightforward way to do so.

Given a set of beat intervals, the union of which makes up the entire song, we take blocks to be all contiguous groups of  $B$  beat intervals. In other words, we create a sequence of overlapping blocks  $X_1, X_2, \dots$  such that  $X_i$  is made up of  $B$  time-contiguous beat intervals, and  $X_i$  and  $X_{i+1}$  differ only by the starting beat of  $X_i$  and the finishing beat of  $X_{i+1}$ . Hence, given  $N$  beat intervals, there are  $N - B + 1$  blocks total. Note that computing an embedding over more than one beat is similar in spirit to the chroma delay embedding approach in [25]. Intuitively, examining patterns over a group of beats gives more information than one beat alone, the effect of which is empirically evaluated in Section 5. For all blocks, we take the window size  $W$  to be the length of the average tempo period, and we advance the window intervals evenly from the beginning of the block to the end of a block with a *hop size*  $H = W/200$ . Hence, there is a 99.5% overlap between windows. We were inspired by theory on raw 1D time series signals [18], which shows that matching the window length to be just under the length of the period in a delay embedding maximizes the roundness of the embedding. Here we would like to match beat-level periodicities and fluctuations therein, so it is sensible to choose a window size corresponding to the tempo. This is in contrast to most other applications that use MFCC sliding window embeddings, which use a much smaller window size on the order of 10s of milliseconds, generally with a 50% overlap, to ensure that the frequency statistics are stationary in each window. In our application, however, we have found that a longer window size makes our self similarity matrices (Section 3.3) smoother, allowing for more reliable matches of beat-level musical trajectories, while having more windows per beat (high overlap) leads to more robust matching of SSMs using L2 (Section 4.1).

Figure 1 shows the first three principal components of an MFCC embedding with a traditional small window size versus our longer window embedding to show the smoothing effect.

### 3.3 Euclidean Self-Similarity Matrices

For each beat-synchronous block  $X_l$  spanning  $B$  beats, we have a 20-dimensional point cloud extracted from the sliding window MFCC representation. Given such a time-ordered point cloud, there is a natural way to create an image which represents the shape of this point cloud in a rotation and translation invariant way, called the *self-similarity matrix* (SSM) representation.

**Definition 1.** A Euclidean Self-Similarity Matrix (SSM) over an ordered point cloud  $X_l \in \mathbb{R}^{M \times k}$  is an  $M \times M$  matrix  $D$  so that

$$D_{ij} = \|X_l[i] - X_l[j]\|_2 \quad (3)$$

In other words, an SSM is an image representing all pairwise distances between points in a point cloud ordered by time. SSMs have been used extensively in the MIR community already, spearheaded by the work of Foote in 2000 for note segmentation in time [10]. They are now often used in general segmentation tasks [24] [15]. They have also been successfully applied in other communities, such as computer vision to




Figure 1: A screenshot from our GUI showing PCA on the sliding window representation of an 8-beat block from the hook of Robert Palmer’s “Addicted To Love” with two different window sizes. Cool colors indicate windows towards the beginning of the block, and hot colors indicate windows towards the end.

recognize activity classes in videos from different points of view and by different actors [13]. Inspired by this work, we use self-similarity matrices as isometry invariant descriptors of local shape in our sliding windows of beat blocks, with the goal of capturing relative shape. In our case, the “activities” are musical expressions over small intervals, and the “actors” are different performers or groups of instruments.

To help normalize for loudness and other changes in relationships between instruments, we first center the point cloud within each block on its mean and scale each point to have unit norm before computing the SSM. That is, we compute the SSM on  $\hat{X}^l$ , where

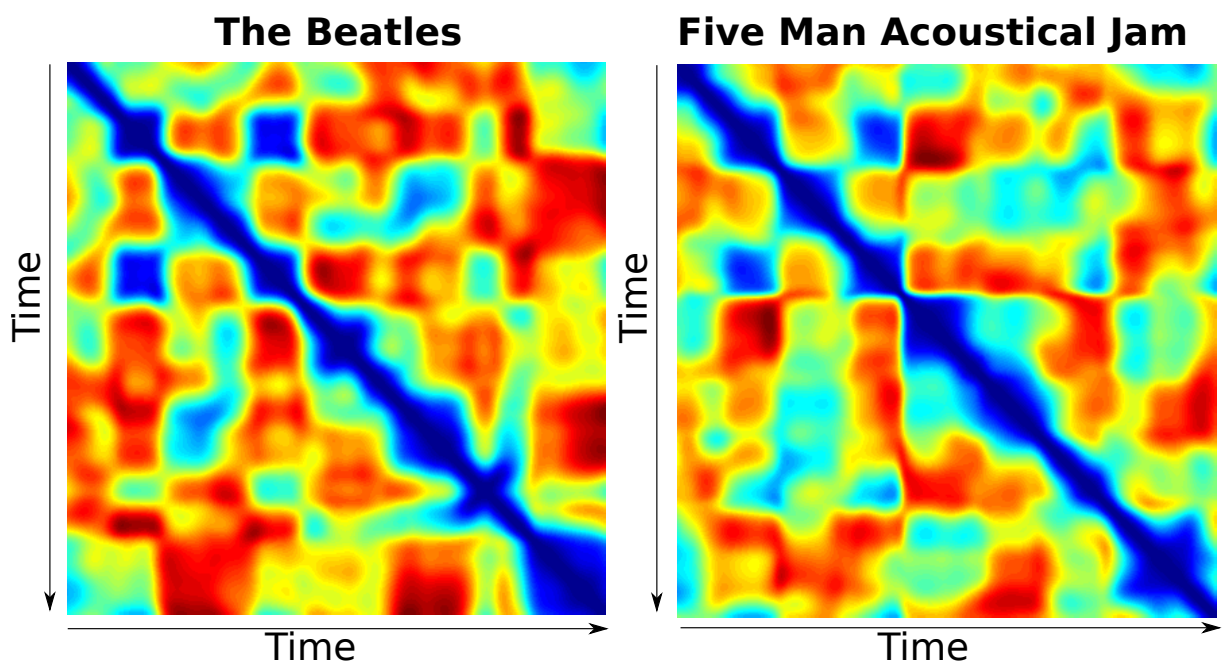
$$\hat{X}_l = \left\{ \frac{x - \text{mean}(x)}{\|x - \text{mean}(x)\|_2} : x \in X_l \right\} \quad (4)$$

Also, not every beat block has the same number of samples due to natural variations of tempo in real songs. Thus, to allow comparisons between all blocks, we resize each SSM to a common image dimension  $d \times d$ , which is a parameter chosen in advance, the effects of which are explored empirically in Section 5.


Figure 2 shows examples of SSMs of 4-beat blocks pulled from the Covers80 dataset that our algorithm matches between two different versions of the same song. Visually, similarities in the matched regions are evident. In particular, viewing the images as height functions, many of the critical points are close to each other. The “We Can Work It Out” example shows how this can work even for live performances, where the overall acoustics are quite different. Even more strikingly, the “Don’t Let It Bring You Down” example shows how similar shape patterns emerge even with an opposite gender singer and radically different instrumentation. Of course, in both examples, there are subtle differences due to embellishments, local time stretching, and imperfect normalization between the different versions, but as we show in Section 5, there are often enough similarities to match up blocks correctly in practice.

## 4 Global Comparison of Two Songs

Once all of the beat-synchronous SSMs have been extracted from two songs, we do a global comparison between all SSMs from two songs to score them as cover matches. Figure 3 shows a block diagram of our system. After extracting beat-synchronous timbral shape features on SSMs, we then extract a binary cross-similarity matrix based on the L2 distance between all pairs of self-similarity matrices between two songs. We subsequently apply the Smith Waterman algorithm on the binary cross-similarity matrix to score a match between the two songs.



(a) A block of 4 beats with 400 windows sliding in the song “We Can Work It Out” by The Beatles with a cover by Five Man Acoustical Jam



(b) A block of 4 beats with 400 windows sliding in the song “Don’t Let It Bring You Down” by Neil Young with a cover by Annie Lennox.

Figure 2: Two examples of MFCC SSM blocks which were matched between a song and its cover in the covers80 dataset. Hot colors indicate windows in the block are far from each other, and cool colors indicate that they are close.




Figure 3: A block diagram of our system for computing a cover song similarity score of two songs using timbral features.

#### 4.1 Binary Cross-Similarity And Local Alignment

Given a set of  $N$  beat-synchronous block SSMs for a song A and a set of  $M$  beat-synchronous block SSMs for a song B, we compute a song-level matching between song A and B by comparing all pairs of SSMs between the two songs. For this we create an  $N \times M$  cross-similarity matrix (CSM), where

$$CSM_{ij} = ||SSMA_i - SSMB_j||_2 \quad (5)$$

is the Frobenius norm (L2 image norm) between the SSM for the  $i^{\text{th}}$  beat block from song A and the SSM for  $j^{\text{th}}$  beat block for song B. Given this cross-similarity information, we then compute a binary cross similarity matrix  $B^M$ . A binary matrix is necessary so that we can apply the Smith Waterman local alignment algorithm [27] to score the match between song A and B, since Smith Waterman only works on a discrete, quantized alphabet, not real values [23]. To compute  $B^M$ , we take the mutual fraction  $\kappa$  nearest neighbors between song A and song B, as in [25]. That is,  $B_{ij}^M = 1$  if  $CSM_{ij}$  is within the  $\kappa M^{\text{th}}$  smallest values in row  $i$  of the CSM and if  $CSM_{ij}$  is within the  $\kappa N^{\text{th}}$  smallest values in column  $j$  of the CSM, and 0 otherwise. As in [25], we found that a dynamic distance threshold for mutual nearest neighbors per element worked significantly better than a fixed distance threshold for the entire matrix.


Once we have the  $B^M$  matrix, we can feed it to the Smith Waterman algorithm, which finds the best local alignment between the two songs, allowing for time shifting and gaps. Local alignment is a more appropriate choice than global alignment for the cover songs problem, since it is possible that different versions of the same song may have intros, outros, or bridge sections that were not present in the original song, but otherwise there are many sections in common. We choose a version of Smith Waterman with diagonal constraints, which was shown to work well for aligning binary cross-similarity matrices for chroma in cover song identification [23]. In particular, we recursively compute a matrix  $D$  so that

$$D_{ij} = \max \left\{ \begin{array}{l} D_{i-1,j-1} + (2\delta(B_{i-1,j-1}) - 1) + \Delta(B_{i-2,j-2}, B_{i-1,j-1}), \\ D_{i-2,j-1} + (2\delta(B_{i-1,j-1}) - 1) + \Delta(B_{i-3,j-2}, B_{i-1,j-1}), \\ D_{i-1,j-2} + (2\delta(B_{i-1,j-1}) - 1) + \Delta(B_{i-2,j-3}, B_{i-1,j-1}), \\ 0 \end{array} \right. \quad (6)$$


where  $\delta$  is the Kronecker delta function and

$$\Delta(a, b) = \left\{ \begin{array}{ll} 0 & b = 1 \\ -0.5 & b = 0, a = 1 \\ -0.7 & b = 0, a = 0 \end{array} \right\} \quad (7)$$


The  $(2\delta(B_{i-1,j-1}) - 1)$  term in each line is such that there will be a +1 score for a match and a -1 score for a mismatch. The  $\Delta$  function is the so-called ‘‘affine gap penalty’’ which gives a score of  $-0.5 - 0.7(g - 1)$  for a gap of length  $g$ . The local constraints are to bias Smith Waterman to choosing paths along near-diagonals of  $B^M$ . This is important since in musical applications, we do not expect large gaps in time in one song that are not in the other, which would show up as horizontal or vertical paths through the  $B^M$  matrix. Rather, we prefer gaps that occur nearly simultaneously in time for a poorly matched beat or set of beats in



(a) Full cross-similarity matrix (CSM)




(b)  $212 \times 212$  Binary cross-similarity matrix ( $B^M$ ) with  $\kappa = 0.05$




(c) Smith Waterman with local constraints: Score 93.1

Figure 4: Cross-similarity matrix and Smith Waterman on MFCC-based SSMS for a true cover song pair of “We Can Work It Out” by The Beatles and Five Man Acoustical Jam.



(a) Full cross-similarity matrix (CSM)

(b)  $212 \times 185$  Binary cross-similarity matrix ( $B^M$ ) with  $\kappa = 0.05$



(c) Smith Waterman with local constraints: Score 8

Figure 5: Cross-similarity matrix and Smith Waterman on MFCC-based SSMs for two songs that are not covers of each other: “We Can Work It Out” by The Beatles and “Yesterday” by En Vogue.




Figure 6: Constrained local matching paths considered in Smith Waterman, as prescribed by [23].



an otherwise well-matching section. Figure 6 shows a visual representation of the paths considered through  $B^M$ .

Figure 4 shows an example of a CSM,  $B^M$ , and resulting Smith Waterman for a true cover song pair. Several long diagonals are visible, indicating large chunks of the two songs are in correspondence, and this gives rise to a large score of 93.1 between the two songs. Figure 5 shows the CSM,  $B$ , and Smith Waterman for two songs which are not versions of each other. By contrast, there are no long diagonals, and this pair only receives a score of 8.

## 5 Results

### 5.1 Covers 80

To benchmark our algorithm, we apply it to the standard ‘‘Covers 80’’ dataset [7], which consists of 80 sets of two versions of the same song, most of which are pop songs from the past three decades. There are designated two sets of songs A and B, each with exactly one version of every pair. To benchmark our algorithm on this dataset, we follow the scheme in [5] and [8]. That is, given a song from set A, compute the Smith Waterman score from all songs from set B and declare the cover song to be the one with the maximum score. Note that a random classifier would only get 1/80 in this scheme. The best scores reported on this dataset are 72/80 [20], using a support vector machine on several different chroma-derived features.

Table 1 shows the correctly identified songs based on the maximum score, given variations of the parameters we have in our algorithm. We achieve a maximum score of 42/80 for a variety of parameter combinations. The nearest neighbor fraction  $\kappa$  and the dimension of the SSM image have very little effect, but increasing the number of beats per block has a positive effect on the performance. The stability of  $\kappa$  and  $d$  are encouraging from a robustness standpoint, and the positive effect increasing the number of beats per block suggests that the shape of medium scale musical expressions are more discriminative than smaller ones.

Table 1: The number of songs that are correctly ranked as the most similar in the Covers 80 dataset, varying parameters.  $\kappa$  is the nearest neighbor fraction,  $B$  is the number of beats per block, and  $d$  is the resized dimension of the Euclidean Self-Similarity images.

| Kappa = 0.05 | B = 8 | B = 10 | B = 12 | B = 14 |
|--------------|-------|--------|--------|--------|
| d = 100      | 30    | 33     | 36     | 40     |
| d = 200      | 31    | 33     | 36     | 39     |
| d = 300      | 31    | 34     | 36     | 40     |
| Kappa = 0.1  | B = 8 | B = 10 | B = 12 | B = 14 |
| d = 100      | 35    | 39     | 41     | 42     |
| d = 200      | 36    | 38     | 42     | 42     |
| d = 300      | 36    | 38     | 41     | 41     |
| Kappa = 0.15 | B = 8 | B = 10 | B = 12 | B = 14 |
| d = 100      | 36    | 42     | 41     | 42     |
| d = 200      | 36    | 41     | 41     | 42     |
| d = 300      | 38    | 42     | 42     | 41     |

In addition to the Covers 80 benchmark, we apply our cover songs score to a recent popular music controversy, the ‘‘Blurred Lines’’ controversy [16]. Marvin Gaye’s estate argues that Robin Thicke’s recent pop song ‘‘Blurred Lines’’ is a copyright infringement of Gaye’s ‘‘Got To Give It Up.’’ Though the note sequences differ between the two songs, ruling out any chance of a high chroma-based score, Robin Thicke has said that his song was meant to ‘‘evoke an era’’ (Marvin Gaye’s era) and that he derived significant inspiration from ‘‘Got To Give It Up’’ specifically [16]. Without making a statement about any legal implications, we note that our timbral shape-based score between ‘‘Blurred Lines’’ and ‘‘Got To Give It Up’’ is in the 99.9<sup>th</sup> percentile of all scores between songs in group A and group B in the Covers 80 dataset, for  $\kappa = 0.1$ ,  $B = 14$ , and  $d = 200$ . Unsurprisingly, when comparing ‘‘Blurred Lines’’ with all other songs in the Covers 80 database plus ‘‘Got To Give It Up,’’ ‘‘Got To Give It Up’’ was the highest ranked. For reference, binary cross similarity matrices are shown in Figure 7, both for our timbre shape based technique and the delay embedding chroma




Figure 7: Corresponding portions of the binary cross-similarity matrix between Marvin Gaye’s “Got To Give It Up” and Robin Thicke’s “Blurred Lines” for both shape-based timbre (our technique) and chroma delay embedding

technique in [25]. The timbre-based cross-similarity matrix is densely populated with diagonals, while the pitch-based one is not.

## 6 Conclusions And Future Work

We show that timbral information in the form of MFCC can indeed be used for cover song identification. Most prior approaches have used Chroma-based features averaged over intervals. By contrast, we show that an analysis of the fine relative shape of MFCC features over intervals is another way to achieve good performance. This opens up the possibility for MFCC to be used in much more flexible music information retrieval scenarios than traditional audio fingerprinting.

On the more technical side, we should note that for comparing shape, L2 of SSMs for cross-similarity is fairly simple and not robust to local re-parameterizations in time between versions, though we tried many other isometry invariant shape descriptors that were significantly slower and yielded inferior performance in initial implementation. In particular, we tried curvature descriptors (ratio of arc length to chord length), Gromov-Hausdorff distance after fractional iterative closest points aligning MFCC block curves [19], and Earth Mover’s distance between SSMs [26]. If we are able to find another shape descriptor which performs better than our current scheme but is slower, we may still be able to make it computationally feasible by using the “Generalized Patch Match” algorithm [1] to reduce the number of pairwise block comparisons needed by exploiting coherence in time. This is similar in spirit to the approximate nearest neighbors schemes proposed in [28] for large scale cover song identification, and we could adapt their sparse Smith Waterman algorithm to our problem. In an initial implementation of generalized patch match for our current scheme, we found we only needed to query about 15% of the block pairs.

## 7 Supplementary Material

We have documented our code and uploaded directions for performing all experiments run in this paper. We also created an open source graphical user interface which can be used to interactively view cross-similarity matrices and to examine the shape of blocks of audio after 3D PCA using OpenGL. All code can be found in the ISMIR2015 directory at

[github.com/ctralie/PublicationsCode](https://github.com/ctralie/PublicationsCode).

## 8 Acknowledgements

Chris Tralie was supported under NSF-DMS 1045133 and an NSF Graduate Fellowship. Paul Bendich was supported by NSF 144749. John Harer and Guillermo Sapiro are thanked for valuable feedback. The authors would also like to thank the Information Initiative at Duke (iiD) for stimulating this collaboration.

## References

- [1] Connelly Barnes, Eli Shechtman, Dan B Goldman, and Adam Finkelstein. The generalized patchmatch correspondence algorithm. In *Computer Vision–ECCV 2010*, pages 29–43. Springer, 2010.
- [2] Juan Pablo Bello. Audio-based cover song retrieval using approximate chord sequences: Testing shifts, gaps, swaps and beats. In *ISMIR*, volume 7, pages 239–244, 2007.
- [3] Bruce P Bogert, Michael JR Healy, and John W Tukey. The quefrency alalysis of time series for echoes: Cepstrum, pseudo-autocovariance, cross-cepstrum and saphé cracking. In *Proceedings of the symposium on time series analysis*, volume 15, pages 209–243. chapter, 1963.
- [4] Michael Casey and Malcolm Slaney. The importance of sequences in musical similarity. In *Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on*, volume 5, pages V–V. IEEE, 2006.
- [5] Daniel PW Ellis. Identifying ‘cover songs’ with beat-synchronous chroma features. *MIREX 2006*, pages 1–4, 2006.
- [6] Daniel PW Ellis. Beat tracking by dynamic programming. *Journal of New Music Research*, 36(1):51–60, 2007.
- [7] Daniel PW Ellis. The “covers80” cover song data set. URL: <http://labrosa.ee.columbia.edu/projects/coversongs/covers80>, 2007.
- [8] Daniel PW Ellis and Courtenay Valentine Cotton. The 2007 labrosa cover song detection system. *MIREX 2007*, 2007.
- [9] Daniel PW Ellis and Bertin-Mahieux Thierry. Large-scale cover song recognition using the 2d fourier transform magnitude. In *The 13th international society for music information retrieval conference*, pages 241–246, 2012.
- [10] Jonathan Foote. Automatic audio segmentation using a measure of audio novelty. In *Multimedia and Expo, 2000. ICME 2000. 2000 IEEE International Conference on*, volume 1, pages 452–455. IEEE, 2000.
- [11] Rémi Foucard, J-L Durrieu, Mathieu Lagrange, and Gaël Richard. Multimodal similarity between musical streams for cover version detection. In *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*, pages 5514–5517. IEEE, 2010.
- [12] Eric J Humphrey, Oriol Nieto, and Juan Pablo Bello. Data driven and discriminative projections for large-scale cover song identification. In *ISMIR*, pages 149–154, 2013.
- [13] Imran N Junejo, Emilie Dexter, Ivan Laptev, and Patrick Pérez. Cross-view action recognition from temporal self-similarities. In *Proceedings of the 10th European Conference on Computer Vision: Part II*, pages 293–306. Springer-Verlag, 2008.
- [14] Samuel Kim, Erdem Unal, and Shrikanth Narayanan. Music fingerprint extraction for classical music cover song identification. In *Multimedia and Expo, 2008 IEEE International Conference on*, pages 1261–1264. IEEE, 2008.
- [15] Brian McFee and Daniel PW Ellis. Analyzing song structure with spectral clustering. In *15th International Society for Music Information Retrieval (ISMIR) Conference*, 2014.

- [16] Emily Miao and Nicole E Grimm. The blurred lines of what constitutes copyright infringement of music: Robin thicke v. marvin gayes estate. *WESTLAW J. INTELLECTUAL PROP.*, 20:1, 2013.
- [17] Oriol Nieto and Juan Pablo Bello. Music segment similarity using 2d-fourier magnitude coefficients. In *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*, pages 664–668. IEEE, 2014.
- [18] Jose A Perea and John Harer. Sliding windows and persistence: An application of topological methods to signal analysis. *Foundations of Computational Mathematics*, pages 1–40, 2013.
- [19] Jeff M Phillips, Ran Liu, and Carlo Tomasi. Outlier robust icp for minimizing fractional rmsd. In *3-D Digital Imaging and Modeling, 2007. 3DIM'07. Sixth International Conference on*, pages 427–434. IEEE, 2007.
- [20] Suman Ravuri and Daniel PW Ellis. Cover song detection: from high scores to general classification. In *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*, pages 65–68. IEEE, 2010.
- [21] Justin Salamon, Joan Serrà, and Emilia Gómez. Melody, bass line, and harmony representations for music version identification. In *Proceedings of the 21st international conference companion on World Wide Web*, pages 887–894. ACM, 2012.
- [22] J Serra. Music similarity based on sequences of descriptors: tonal features applied to audio cover song identification. *Department of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona, Spain*, 2007.
- [23] Joan Serra, Emilia Gómez, Perfecto Herrera, and Xavier Serra. Chroma binary similarity and local alignment applied to cover song identification. *Audio, Speech, and Language Processing, IEEE Transactions on*, 16(6):1138–1151, 2008.
- [24] Joan Serra, Meinard Müller, Peter Grosche, and Josep Lluís Arcos. Unsupervised detection of music boundaries by time series structure features. In *Twenty-Sixth AAAI Conference on Artificial Intelligence*, 2012.
- [25] Joan Serra, Xavier Serra, and Ralph G Andrzejak. Cross recurrence quantification for cover song identification. *New Journal of Physics*, 11(9):093017, 2009.
- [26] Sameer Shirdhonkar and David W Jacobs. Approximate earth movers distance in linear time. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.
- [27] Temple F Smith and Michael S Waterman. Identification of common molecular subsequences. *Journal of molecular biology*, 147(1):195–197, 1981.
- [28] Romain Tavenard, Hervé Jégou, and Mathieu Lagrange. Efficient cover song identification using approximate nearest neighbors. 2012.
- [29] Julián Urbano, Juan Lloréns, Jorge Morato, and Sonia Sánchez-Cuadrado. Melodic similarity through shape similarity. In *Exploring music contents*, pages 338–355. Springer, 2011.