



HAL
open science

Barwise Music Structure Analysis with the Correlation Block-Matching Segmentation Algorithm

Axel Marmoret, Jérémy E Cohen, Frédéric Bimbot

► **To cite this version:**

Axel Marmoret, Jérémy E Cohen, Frédéric Bimbot. Barwise Music Structure Analysis with the Correlation Block-Matching Segmentation Algorithm. Transactions of the International Society for Music Information Retrieval (TISMIR), 2023, 6 (1), pp.167-185. 10.5334/tismir.167 . hal-04323556

HAL Id: hal-04323556

<https://hal.science/hal-04323556v1>

Submitted on 5 Dec 2023

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.



Distributed under a Creative Commons Attribution 4.0 International License



Barwise Music Structure Analysis with the Correlation Block-Matching Segmentation Algorithm

AXEL MARMORET

JÉRÉMY E. COHEN

FRÉDÉRIC BIMBOT

*Author affiliations can be found in the back matter of this article

RESEARCH

]u[ubiquity press

ABSTRACT

Music Structure Analysis (MSA) is a Music Information Retrieval task consisting of representing a song in a simplified, organized manner by breaking it down into sections typically corresponding to “chorus”, “verse”, “solo”, etc. In this work, we extend an MSA algorithm called the Correlation Block-Matching (CBM) algorithm introduced by (Marmoret et al., 2020, 2022b). The CBM algorithm is a dynamic programming algorithm that segments self-similarity matrices, which are a standard description used in MSA and in numerous other applications. In this work, self-similarity matrices are computed from the feature representation of an audio signal and time is sampled at the bar-scale. This study examines three different standard similarity functions for the computation of self-similarity matrices. Results show that, in optimal conditions, the proposed algorithm achieves a level of performance which is competitive with supervised state-of-the-art methods while only requiring knowledge of bar positions. In addition, the algorithm is made open-source and is highly customizable.

CORRESPONDING AUTHOR:
Axel Marmoret

Univ. Rennes 1, Inria, CNRS, IRISA, France; IMT Atlantique, Lab-STICC, Brest, France
axel.marmoret@imt-atlantique.fr

KEYWORDS:

Music Structure Analysis; Audio Signals; Barwise Music Processing; Self-Similarity Matrix Segmentation

TO CITE THIS ARTICLE:

Marmoret, A., Cohen, J. E., and Bimbot, F. (2023). Barwise Music Structure Analysis with the Correlation Block-Matching Segmentation Algorithm. *Transactions of the International Society for Music Information Retrieval*, 6(1), 167–185. DOI: <https://doi.org/10.5334/tismir.167>

1. INTRODUCTION

Citing Paulus et al. (2010), “[...] it is the structure, or the relationships between the sound events that create musical meaning”. In that sense, researchers in MIR developed the Music Structure Analysis (MSA) task, which focuses on the retrieval of the *structure* in a song. Music structure is ill-defined, but is generally viewed as a hierarchical description, from the level of notes to the level of the song itself (McFee et al., 2017; Nieto et al., 2020). A tentative definition is that structure is *a simplified representation of the organization of the song*.

In that sense, motifs which arise from the organization of notes are a first level of structure. These motifs create patterns, progressions and phrases. In general, the highest level of structure defines musical sections, corresponding to “chorus”, “verse” and “solo”, which is a macroscopic description of music (Sargent et al., 2016). Some work focuses on estimating structure in its hierarchical nature (e.g. McFee and Ellis, 2014a, b; de Berardinis et al., 2020; Salamon et al., 2021), but this work focuses on a “flat” level of segmentation, *i.e.* a macroscopic level, corresponding to musical sections. Facing the high diversity of music, and the many ways structure can be designed, we restrict this work to the study of Western modern (and in particular Western Popular) music. In particular, this work relies on both the RWC Pop (Goto et al., 2002) and the SALAMI (Smith et al., 2011) datasets, which are open-source and standard datasets in MSA.

MSA is subdivided into two subtasks, not necessarily mutually exclusive: the **boundary retrieval** task and the **segment labelling** task. The boundary retrieval task consists in estimating the boundaries between different sections, hence partitioning music into several non-overlapping segments, covering the entire song. The segment labelling task consists in grouping similar segments with the same label, typically letters such as ‘A’, ‘B’, ‘C’, etc. In this article, only the boundary retrieval task is considered. A schematic example of musical structure is presented in Figure 1.

1.1 RELATED WORK

Algorithms aimed at solving MSA are designed according to one or several criteria among the following:

homogeneity, novelty, repetition and **regularity** (Nieto et al., 2020). The homogeneity criterion assumes that a section consists of similar musical elements (notes, chords, tonality, timbre, ...). Novelty is the counterpart of homogeneity: this criterion considers that boundaries are located primarily between consecutive musical elements that are highly dissimilar. A high novelty is salient between two distinct homogeneous zones (“break” of homogeneity), and conversely, homogeneity is evaluated within successive dissimilar portions in the song. The third criterion, repetition, relies on a global approach to the song. The rationale is that a motif (e.g. a melodic line) may be time-varying (and thus heterogeneous), but can define a segment if it is repeated across the song (for example, a chorus). The repetition criterion may also be used to partition long segments into smaller repeated segments. Finally, the regularity criterion assumes that, within a song (and even within a musical genre), segments should be of comparable size.

Many MSA algorithms make use of matrices representing the similarity and dissimilarity in music, sometimes referred to as “self-distance matrices” (Paulus et al., 2010), “self-similarity matrices” (Nieto et al., 2020), “recurrence matrices” or “pair-wise frame similarities” (McFee and Ellis, 2014a). These representations differ in their details, but share the same conceptual idea of computing some form of similarity (or, conversely, some form of distance) between the different frames of music, and representing it in a square matrix (its size being the number of frames). In this work, we will use the term “self-similarity matrix”.

As a particular example, the novelty kernel (Foote, 2000), which may be some of the earliest work on audio MSA, estimates boundaries as points of high dissimilarity between the recent past and the near future, by applying a square kernel matrix on the diagonal of the self-similarity matrix. This kernel works ideally when the recent past and the near future are homogenous in their respective neighborhoods, but very dissimilar with each other. In practice, this kernel is convolved with the self-similarity matrix of the song, centered on the diagonal, which gives a “novelty” value for each temporal sample of the song, finally post-processed into boundaries with a thresholding operation.

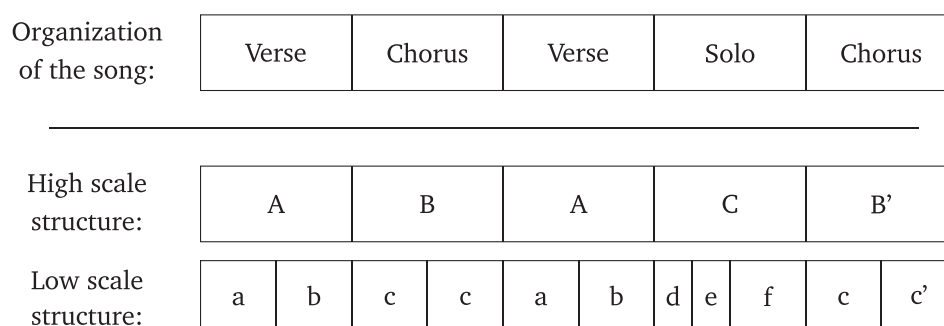


Figure 1 A schematic example of musical structure.

While this technique is rather simple, it is still used as a standard segmentation tool in recent work (e.g. McCallum, 2019; Wang et al., 2021) focusing on improving the boundary retrieval performance by enhancing the self-similarity matrix. In particular, both of these works belong to the domain of representation learning (Bengio et al., 2013), consisting of designing machine learning algorithms to learn relevant representations instead of focusing on solving a particular task. In that context, using prior knowledge, both McCallum (2019) and Wang et al. (2021) design neural network architectures and optimization schemes with the objective to obtain enhanced (nonlinear) similarity functions, more prone to highlight the structure in the self-similarity matrices.

Notably, McCallum (2019) develops an unsupervised learning scheme where the prior knowledge enforced in the representation is based on the proximity of samples: the closer the frames in the song, the more probable they belong to the same segment. In the same spirit, Wang et al. (2021) develop a supervised learning scheme: the neural network learns representations where segments annotated with the same label are close, and segments annotated differently are far apart. The rationale for both methods is to learn a similarity function which is not only representing the feature-wise correspondence of two music frames, but can also discover frequent patterns in the learning samples.

McFee and Ellis (2014a) propose an algorithm based on spectral clustering, aiming at interpreting the repetitive patterns in a song as principally connected vertices in a graph. The structure is then obtained by studying the eigenvectors of the Laplacian of this graph, forming cluster classes for segmentation. This technique is amongst the best-performing unsupervised techniques nowadays, and was improved by recent work by Salamon et al. (2021), which replaces or enhances the acoustic features on which spectral clustering is applied with nonlinear embeddings, learned by means of a neural network.

Serrà et al. (2014) develop “Structural Features”, which, by design, encode both repetitive and homogeneous parts. The rationale of these features is to compute the similarity between bags of instances, composed of several consecutive frames. In that sense, the similarity encodes the repetition of any sequence, which can be stationary (homogeneity) or varying (repetition). Boundaries are obtained as points of high novelty between consecutive structural features.

Finally, Grill and Schlüter (2015) develop a Convolutional Neural Network (CNN) which outputs estimated boundaries. This CNN is one of the few techniques which does not compute a self-similarity matrix to later post-process it into boundaries, but it still uses self-similarities as input. The network is supervised on two-level annotations, on the SALAMI dataset (Smith

et al., 2011), and, according to the authors, using these two levels of annotations is beneficial to the performance.

While many algorithms are devoted to the task of boundary retrieval (see for instance literature reviews from Paulus et al. (2010) and from Nieto et al. (2020)), research is still conducted towards more effective estimation algorithms. As presented above, in the past decade, research has mainly shifted from unsupervised to supervised algorithms, *i.e.* from low-informed estimation algorithms, generally designed with strong hypotheses, to algorithms which take advantage of (generally huge) annotated databases to learn mappings between the musical features and annotated structural elements. While this shift has resulted in more effective algorithms, it has the disadvantages of requiring large training datasets, and reproducing potential bias in relation to the annotations, known to be prone to high subjectivity and ambiguity (Nieto et al., 2020).

1.2 CONTRIBUTIONS

In order to improve unsupervised algorithms, we propose in this article a novel approach based on the Correlation “Block-Matching” (CBM) algorithm. This algorithm was briefly introduced in previous work (Marmoret et al., 2020, 2022b) and is worth a more detailed presentation, which is one of the objectives of this work. Firstly, in line with the findings in (Marmoret et al., 2020, 2022b), we conjecture that the bar-scale is the most appropriate temporal scale from which to infer structure in Western modern music, and we present a framework which inherently processes music at this temporal scale. To the best of our knowledge, only a few works used such a hypothesis (e.g. Wang et al., 2021; Marmoret et al., 2020, 2022b). This hypothesis is supported by experiments which compare segmentation performance when aligning state-of-the-art algorithms and the CBM algorithm on either the beat or the bar-scale. We show a consistent advantage for the bar-scale alignment approach.

The CBM algorithm estimates the musical structure based on the principles introduced in the work of Jensen (2006), later extended by Sargent et al. (2016). In a nutshell, the CBM algorithm is based on the definition of a score function (further denoted as u) applied to segments, with the overall segmentation of the song resulting in the maximum total score of the set of segments. This defines an optimization problem, which can be solved by dynamic programming.

The novelty of the CBM algorithm lies in its ability to extend previous work by incorporating new hypotheses regarding the design of the score function u . As a consequence, the algorithm is highly customizable and can be tailored to specific hypotheses and applications, which is a potential area for future research. We also present a study of different similarity functions to account for the similarity between musical features, and notably the Radial Basis function, which, to the best

of our knowledge, was never previously used for MSA. Finally, we present experimental results which appear competitive with the most effective algorithm known to date (Grill and Schlüter, 2015).

The CBM algorithm is unsupervised in the sense that the segment boundaries are estimated as solutions of an optimization problem which does not depend explicitly on annotated examples. Nonetheless, in order to accurately tune internal hyperparameters, the following experiments are carried out by separating data between a “train” and a “test” dataset. In addition, we acknowledge that we use a learning-based toolbox for the bar estimation, but this toolbox is independent from our work. In that sense, although we label this algorithm as “unsupervised”, it could also arguably be qualified as “weakly-” or “semi-” supervised.

This article is organized as follows: Section 2 presents in more detail the hypotheses and framework to process music in a barwise setting, Section 3 presents the CBM algorithm and Section 4 presents an evaluation of the CBM algorithm on the boundary retrieval task, along with a comparison with state-of-the-art algorithms.

2. BARWISE MUSIC ANALYSIS

In most work in MSA (Nieto et al., 2020), the signal of a song is represented as time-sampled features, related to some extent to the frequency content of the song’s signal. In previous work on MSA, features have been either computed with a fixed hop length, typically between 0.1s and 1s according to Paulus et al. (2010), or (in more recent work), aligned on beats (McCallum, 2019; Wang et al., 2021; Salamon et al., 2021). Beat alignment is musically-relevant because it aligns the features and the estimations with respect to a time segmentation consistent with music performance. In this work, we hypothesize that the bar-scale is more relevant than the beat-scale to study MSA in Western modern music.

Bars seem well suited to express patterns and sections in Western modern music. Indeed, in Western musical notation, musical note lengths are expressed relatively to beats, and beats are combined to form bars. Bars finally segment the musical scores (with vertical lines), and similarities occur generally across different bars (which is particularly visible by the use of repeat bars, or symbols as “Dal Segno”, “Da Capo”, etc). In addition, the intuition that musical sections are synchronized on downbeats is experimentally confirmed by Mauch et al. (2009) and Fuentes et al. (2019), where the use of structural information improves the estimation of downbeats. Experiments supporting this hypothesis are presented in Section 2.3.

The direct drawback of barwise alignment is the need for a powerful tool to estimate bar boundaries. In this work, we use the *madmom* toolbox (Böck et al., 2016a),

which uses a neural network to perform bar estimation (Böck et al., 2016b). In the 2016 MIREX contest,¹ which was the last edition of the contest comparing downbeat estimation algorithms, this neural network obtained the best performance, and can hence be considered as one of the state-of-the-art algorithms for the task. Even if some algorithms obtained better performance since (e.g. Böck and Davies, 2020; Oyama et al., 2021; Hung et al., 2022), we consider that the *madmom* toolbox achieves a satisfactory level of performance for our intended application.

2.1 BARWISE TF MATRIX

In this work, we represent music as barwise spectrograms, and more particularly as a **Barwise TF matrix**, following Marmoret et al. (2022b). The Barwise TF matrix consists of a matrix of size $B \times TF$, B being the number of bars in the song (i.e. a dimension accounting for the bar-scale), and TF the vectorization of both time (at bar-scale) and feature dimensions (representing the frequency to some extent) into a unique Time-Frequency dimension. The number of time frames per bar is fixed to $T = 96$, as in Marmoret et al. (2022b). Following the work of Grill and Schlüter (2015), the signal is represented in log mel features, i.e. the logarithm of mel coefficients, expressed with $F = 80$ mel coefficients, but any other feature representation could be used instead. The rationale for using log mel spectrograms is that they lead to high segmentation performance (Grill and Schlüter, 2015; Nieto et al., 2020) while constituting a compact spectral representation, suited for music analysis.

2.2 BARWISE SELF-SIMILARITY MATRIX

As stated in Section 1.1, a common representation in MSA is the self-similarity matrix, representing the similarities at the scale of the song. An idealized self-similarity matrix, extracted from Paulus et al. (2010), is presented in Figure 2.

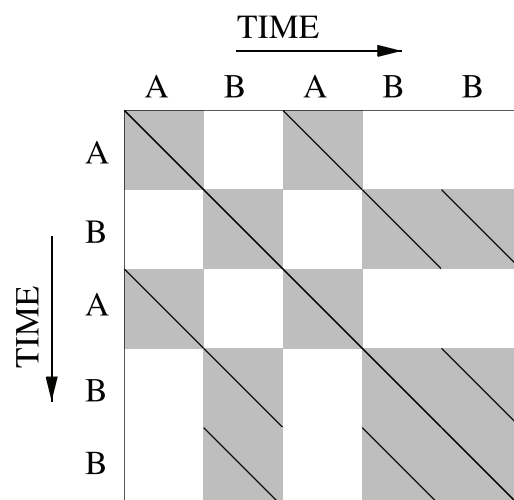


Figure 2 An idealized self-similarity matrix, extracted from Paulus et al. (2010).

Similar passages are identified by two typical shapes: **blocks** and **stripes**. A block is a square (or a rectangle) in the self-similarity matrix, representing a zone of high inner-similarity, *i.e.* several consecutive frames which are highly similar, hence corresponding to the homogeneity criterion. A stripe is a line parallel to the main diagonal representing a repetition of the content, *i.e.* a pattern of several frames repeated in the same order, hence corresponding to the repetition criterion. As a general trend, the segmentation algorithms using self-similarity matrices are designed so as to retrieve segments based on blocks and stripes.

Given a Barwise TF matrix $X \in \mathbb{R}^{B \times TF}$, the self-similarity matrix of X is defined as $A(X) \in \mathbb{R}^{B \times B}$ where each coefficient (i, j) represents the similarity between vectors $X_i, X_j \in \mathbb{R}^{TF}$. Self-similarity matrices are computed from the Barwise TF representation of the song, therefore, each coefficient in the self-similarity matrix represents the feature-wise similarity for a pair of bars.

The similarity between two vectors is subject to a similarity function (the dot product for instance), and, as a consequence, different self-similarity matrices can be constructed. The main diagonal in a self-similarity matrix represents the self-similarity of each vector, and is in general (and in this work in particular) normalized to one. This work studies three different similarity functions, namely the Cosine, Autocorrelation and RBF similarity functions. The latter two represent novel contributions compared to our previous work (Marmoret et al., 2022b).

2.2.1 Cosine self-similarity matrix

The Cosine similarity function computes the normalized dot product between two vectors, and leads to the Cosine self-similarity matrix, denoted as $A_{\cos}(X)$. Practically, denoting as \tilde{X} the row-wise l_2 -normalized version of X (*i.e.* the matrix X where each row has been divided by its l_2 -norm), the Cosine self-similarity matrix is defined as $A_{\cos}(X) = \tilde{X}\tilde{X}^T$, or, elementwise, for $1 \leq i, j \leq B$:

$$A_{\cos}(X)_{ij} = \frac{\langle X_i, X_j \rangle}{\|X_i\|_2 \|X_j\|_2} = \sum_{k=1}^{TF} \tilde{X}_{ik} \tilde{X}_{jk}. \quad (1)$$

2.2.2 Autocorrelation self-similarity matrix

The Autocorrelation similarity function is defined for 2 bars X_i and X_j as $\text{corr}(X_i, X_j) = (X_i - \bar{x})(X_j - \bar{x})^T$, denoting as $\bar{x} \in \mathbb{R}^{TF}$ the mean of all bars in the song.

The barwise Autocorrelation similarity function yields the Autocorrelation self-similarity matrix $A_{\text{corr}}(X)$ as:

$$A_{\text{corr}}(X)_{ij} = \frac{\langle X_i - \bar{x}, X_j - \bar{x} \rangle}{\|X_i - \bar{x}\|_2 \|X_j - \bar{x}\|_2}. \quad (2)$$

In other words, the Autocorrelation matrix is exactly the Cosine self-similarity matrix of the centered matrix $X - \mathbf{1}_B^T \bar{x}$, *i.e.* $A_{\text{corr}}(X) = A_{\cos}(X - \mathbf{1}_B^T \bar{x})$.

2.2.3 RBF self-similarity matrix

Kernel functions are symmetric positive definite or semi-definite functions. In machine learning, kernel functions are generally used to represent data in a high-dimensional space (sometimes infinite), enabling a nonlinear processing of data with linear methods (*e.g.* nonlinear classification with Support Vector Machines, SVM).

The Radial Basis Function (RBF) kernel is a kernel function defined as $\text{RBF}(X_i, X_j) = \exp(-\gamma \|X_i - X_j\|_2^2)$, γ being a user-defined parameter. The RBF can be used as a similarity function between two bars X_i and X_j , hence defining the RBF self-similarity matrix $A_{\text{RBF}}(X)$ as:

$$A_{\text{RBF}}(X)_{ij} = \text{RBF}(\tilde{X}_i, \tilde{X}_j) = \exp\left(-\gamma \left\| \frac{X_i}{\|X_i\|_2} - \frac{X_j}{\|X_j\|_2} \right\|_2^2\right). \quad (3)$$

Bars are normalized by their l_2 norm in the computation of A_{RBF} , in order to limit the impact of variations of power between bars. The self-similarity of a bar is equal to $e^0 = 1$.

Parameter γ is set relatively to the standard deviation of the pairwise Euclidean distances of all bars in the original matrix (self-distances excluded), to adapt the shape of the exponential function to the relative distribution of distances in the song. Hence, denoting as:

$$\sigma = \text{std}_{\substack{1 \leq i, j \leq B \\ i \neq j}} \left(\left\| \frac{X_i}{\|X_i\|_2} - \frac{X_j}{\|X_j\|_2} \right\|_2 \right), \text{ we set } \gamma = \frac{1}{2\sigma}.$$

The RBF function may be useful for MSA as the self-similarity of dissimilar elements fades rapidly due to the properties of the exponential function. Hence, the RBF similarity function emphasizes similar components (*i.e.* homogeneous zones). The three self-similarities are presented in Figure 3, on the Barwise TF of song POP01 from RWC Pop.

2.3 BARWISE MSA EXPERIMENTS

Section 2 is based on the hypothesis that the bar-scale is more relevant than other time discretizations (in particular the beat-scale) to study MSA in Western modern music. To support this hypothesis, we present hereafter three experiments studying the differences in performance between beat-aligned and barwise-aligned estimations.

2.3.1 Aligning the annotations on the downbeats

As a preliminary experiment to study the impact of barwise alignment on the segmentation quality, we evaluate the loss in performance when aligning annotations on downbeats for both the RWC Pop (Goto et al., 2002) and SALAMI (Smith et al., 2011) datasets. In this experiment, each annotation is aligned with the closest estimated downbeat, and the barwise-aligned annotations are compared with the initial annotation using the standard metrics $P_{0.5s}$, $R_{0.5s}$, $F_{0.5s}$ and P_{3s} , R_{3s} , F_{3s} (detailed in Section 4.1). Results are presented in Table 1. The RWC Pop annotations are barely impacted by the

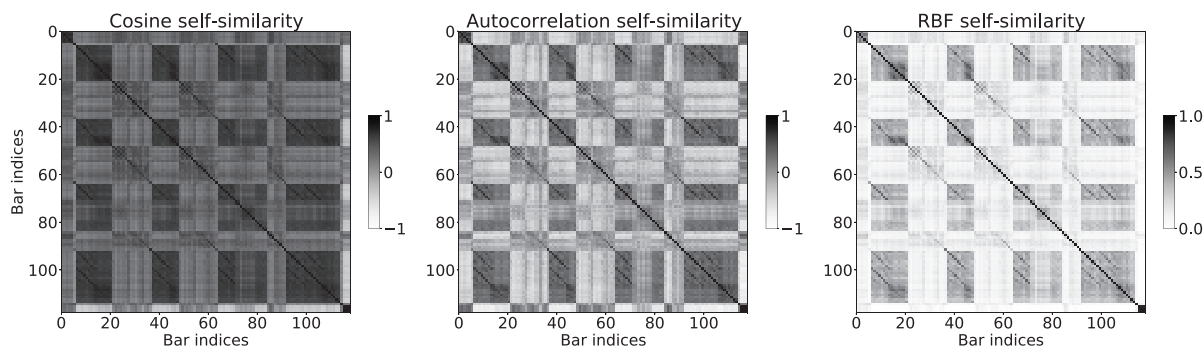


Figure 3 Cosine, Autocorrelation and RBF self-similarities for the song *POP01* of RWC Pop.

Dataset		$P_{0.5s}$	$R_{0.5s}$	$F_{0.5s}$	P_{3s}	R_{3s}	F_{3s}
SALAMI	Annotation 1	82.47%	82.14%	82.30%	99.94%	99.56%	99.74%
	Annotation 2	80.97%	80.92%	80.94%	99.92%	99.84%	99.88%
RWC Pop		96.46%	96.21%	96.33%	100%	99.73%	99.86%

Table 1 Standard metrics (see Section 4.1) when aligning the reference annotations on the downbeats (compared to the original annotations).

barwise alignment, suggesting that annotations are precisely located on downbeats. A loss in performance with short tolerances is observed on the annotations of the SALAMI dataset, either suggesting imprecise bar estimations or boundaries not located on the downbeats. Still, the levels of performance exhibited in Table 1 largely outperform the current state-of-the-art ($\approx 80\%$ vs $\approx 54\%$ for the $F_{0.5s}$ metric, respectively for the downbeat-aligned annotations in Table 1 and for Grill and Schlüter (2015), whose results are presented in Figure 12). In that sense, the loss in performance induced by downbeat alignment may be compensated if estimations are indeed more precise due to this alignment.

2.3.2 Downbeat-alignment for several state-of-the-art algorithms

A second experiment consists of post-processing the boundary estimations of three unsupervised state-of-the-art algorithms (Foote, 2000; McFee and Ellis, 2014a; Serrà et al., 2014), computed with the *MSAF* toolbox (Nieto and Bello, 2016), by aligning each boundary with the closest estimated downbeat. As these algorithms originally use beat-aligned features (resulting in beat-aligned estimations), this experiment compares beat-aligned estimations with downbeat-aligned estimations. Segmentation scores are presented in Figure 4 for both SALAMI and RWC Pop datasets.

Results show that aligning estimated boundaries on downbeats results in a strong increase in performance for $F_{0.5s}$, and to comparable results for F_{3s} , on both datasets. Hence, aligned on downbeats, estimations are more accurate, but the F_{3s} metric is not significantly impacted by this alignment. These results suggest that downbeat-alignment is beneficial on these datasets.

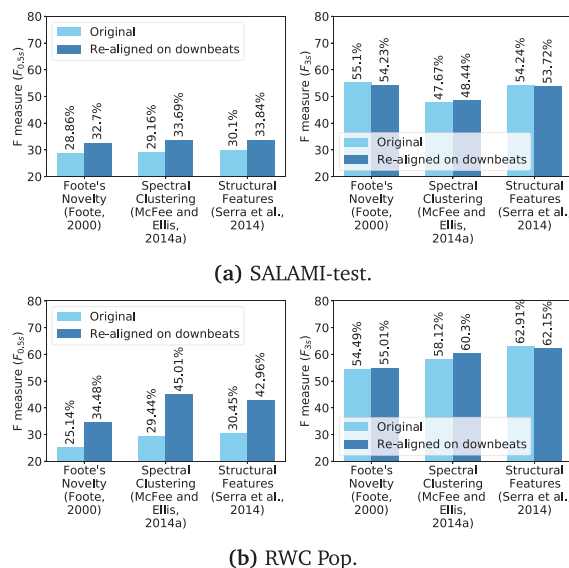


Figure 4 Segmentation results of state-of-the-art algorithms on the SALAMI-test and RWC Pop datasets, for beat-aligned (original) vs. downbeat-aligned boundaries. The SALAMI-test dataset is defined by Ullrich et al. (2014), and introduced in Section 4.2.1.

2.3.3 Focusing on Foote’s algorithm

Finally, a third experiment compares the results obtained with different time discretizations for Foote’s algorithm (Foote, 2000), implemented in the *MSAF* toolbox (Nieto and Bello, 2016). In particular, this experiment compares the results when self-similarities are computed with beat-aligned and downbeat-aligned features.

As presented in Section 1.1, Foote’s algorithm estimates boundaries as points of high novelty. A novelty score is computed at each point of the self-similarity matrix by applying a kernel matrix, and this novelty score is finally post-processed into boundaries with a thresholding operation. In the original implementation,

the Cosine self-similarity is computed on beat-synchronized features, *i.e.* one feature per beat. In this experiment, beats are estimated with the algorithm of Böck et al. (2019), which is one of the state-of-the-art algorithms in beat estimation, and is implemented in the *madmom* toolbox.

As in Section 2.3.2, the original results are compared with the ones obtained when aligning the estimations on downbeats. In addition, we compare the beat-synchronized results with two new feature processing approaches: **bar-synchronized features**, *i.e.* one feature per bar (instead of one per beat), and the **Barwise TF matrix**,² introduced in Section 2.2.

The kernel is of size 66 for the beat-synchronized features (as originally set in MSAF), and is of size 16 for both the bar-synchronized features and the Barwise TF matrix. Results are presented in Tables 2 and 3, respectively for the SALAMI-test and the RWC Pop dataset. In order to fairly compare the algorithms, we also fitted two hyperparameters of the original MSAF implementation for the bar-scale, namely the size of median filtering applied to the input spectrogram and the standard deviation in the Gaussian filter applied to the novelty curve. These parameters were fitted in a train/test fashion, as detailed in Section 4.2.1.

Both Tables 2 and 3 conclude in the same direction: the best performance for the $F_{0.5s}$ and F_{3s} metrics is obtained with the Barwise TF matrix. Similarly to Section 2.3.2, aligning beat-aligned estimations to downbeats in post-processing increases the performance for the $F_{0.5s}$ metric, indicating more precise estimations. It is worthwhile noting that, using bar-synchronized instead of beat-synchronized features increases the performance on the SALAMI dataset, while it decreases it on the RWC Pop dataset.

However, the Barwise TF matrix representation appears to be beneficial for MSA on both datasets.

In future experiments, we denote as “Foote-TF” the condition where Foote’s algorithm is applied to the Barwise-TF matrix, whose results are shown in Tables 2 and 3.

3. CORRELATION “BLOCK-MATCHING” ALGORITHM

With a self-similarity matrix as input, the Correlation “Block-Matching” segmentation algorithm (CBM) estimates boundaries by means of dynamic programming. This algorithm is detailed in this section, along with a study of important parameter settings which were not discussed in the previous work (Marmoret et al., 2020, 2022b), such as the block weighting kernels and the penalty functions. The CBM algorithm estimates boundaries based on the homogeneity/novelty and regularity criteria. The principles of dynamic programming are presented first, followed by the definition of a score function u applied on segments.

3.1 DYNAMIC PROGRAMMING FOR BOUNDARY RETRIEVAL

3.1.1 Boundary retrieval problem

Given a music piece (song) sampled in time as N time steps, the subtask of boundary retrieval can be defined as finding a segmentation (set of boundaries) Z representing the start of each segment, *i.e.* $Z = \{\zeta_i \in \llbracket 1, N \rrbracket, i \in \llbracket 1, E \rrbracket\}$, E representing the number of boundaries estimated in this song. The set of admissible segmentations is denoted as Θ , *i.e.* $Z \in \Theta$. Each segment S_i is composed of the time steps between two consecutive boundaries, *i.e.* $S_i = \{l \in \llbracket 1, N \rrbracket \mid \zeta_i \leq l < \zeta_{i+1}\}$. The second bound is exclusive as it represents the start of the next segment S_{i+1} . By definition, E boundaries define $E-1$ segments. Boundary ζ_i is called the **antecedent** of boundary ζ_{i+1} .

Time synchronization		$P_{0.5s}$	$R_{0.5s}$	$F_{0.5s}$	P_{3s}	R_{3s}	F_{3s}
Beat-synchronized	Original	26.98%	34.58%	29.21%	50.10%	63.30%	54.02%
	Re-aligned on downbeats	31.05%	39.15%	33.33%	50.08%	62.95%	53.78%
Bar-synchronized		37.68%	36.36%	35.97%	58.06%	56.11%	55.57%
Barwise TF Matrix		39.22%	42.66%	39.67%	59.60%	64.82%	60.36%

Table 2 Different time synchronizations for the Foote (2000) algorithm on the SALAMI-test dataset. The SALAMI-test dataset is defined by Ullrich et al. (2014), and introduced in Section 4.2.1.

Time synchronization		$P_{0.5s}$	$R_{0.5s}$	$F_{0.5s}$	P_{3s}	R_{3s}	F_{3s}
Beat-synchronized	Original	31.86%	24.38%	27.29%	67.21%	51.92%	57.95%
	Re-aligned on downbeats	42.30%	32.82%	36.52%	66.67%	51.44%	57.44%
Bar-synchronized		43.53%	26.32%	32.46%	69.25%	42.22%	51.97%
Barwise TF Matrix		53.09%	37.19%	43.30%	79.35%	56.03%	65.04%

Table 3 Different time synchronizations for the Foote (2000) algorithm on the RWC Pop dataset.

3.1.2 Barwise boundary retrieval problem

In the proposed barwise paradigm, the song is discretized into B bars using $N = B + 1$ bar boundaries. Hence, the first boundary is the start of the song, *i.e.* $\zeta_1 = 1$, the last boundary is the end of the last bar in the song,³ *i.e.* $\zeta_E = B + 1$, and each boundary is located on a bar, *i.e.* $\forall i, \zeta_i \in \llbracket 1, B + 1 \rrbracket$. Each segment S_i is composed of the bar indices between two consecutive boundaries.

As a consequence, there exists⁴ $\binom{B-1}{E}$ different sets of boundaries composed of exactly E boundaries, and, more generally, at most $\sum_{k=0}^{B-1} \binom{B-1}{k} = 2^{B-1}$ segmentations for each song. Hence, the segmentation problem admits a finite number of solutions, which can theoretically be solved in a combinatorial way. In practice though, evaluating all possible segmentations leads to an algorithm of exponential complexity $\mathcal{O}(2^B)$, considered intractable in practice.

3.1.3 Dynamic programming

The boundary retrieval problem can be approached as an optimization problem (Jensen, 2006; Sargent et al., 2016). In particular, by associating a score $u(S)$ to each potential segment S , the optimal segmentation Z^* is the segmentation maximizing⁵ the sum of all its segment scores:

$$\begin{aligned} Z^* &= \operatorname{argmax}_{Z \in \Theta} \sum_{i=1}^{E-1} u(\llbracket \zeta_i, \zeta_{i+1} - 1 \rrbracket) \\ &= \operatorname{argmax}_{Z \in \Theta} u(Z) \end{aligned} \quad (4)$$

by extending notation u for a set of segments.

The problem can be solved using a dynamic programming algorithm (Bellman, 1952; Cormen et al., 2009, Chap. 15), the principle of which is to solve a combinatorial optimization problem by dividing it into several independent subproblems. The independent subproblems are formulated in a recursive manner, and their solutions can be stitched together to form a solution to the original problem. Notice that in the current formulation of the segmentation problem, defined in Equation 4, each potential segment is evaluated independently, via its score, and is never compared with the others. In other terms, repetitions of the same section are not considered, while they could inform on the overall structure, typically considering the repetition criterion. Thus, the segmentation problem defined in Equation 4 is a relaxation of the general segmentation problem. This relaxation is considered because it allows to use principles of dynamic programming, by evaluating the score of all segments as independent subproblems. In particular, this relaxed problem is said to exhibit “optimal substructure” (Cormen et al., 2009).

3.1.4 Longest-path on a directed acyclic graph

Following the formulation of Jensen (2006), the segmentation problem can be reframed into the problem of finding the longest path on a Directed Acyclic Graph

(DAG). The rationale of the solution algorithm is that the optimal segmentation up to any given bar b_k can be found exactly by recursively evaluating the optimal segmentations up to each antecedent of b_k , *i.e.* (without any constraint) all bars $b_l < b_k$, and the score of the segments $\llbracket b_l, b_k - 1 \rrbracket$. Formally, denoting as $Z_{\llbracket 1, b_k \rrbracket}^*$ the optimal segmentation up to bar b_k , the CBM algorithm consists of:

1. Lookup for $\{u(Z_{\llbracket 1, b_l \rrbracket}^*), \forall b_l < b_k\}$, *i.e.* the optimal segmentation up to each antecedent, which is stored in an array when first computed,
2. Computing $\{u(\llbracket b_l, b_k - 1 \rrbracket), \forall b_l < b_k\}$, *i.e.* the segmentation score between bars b_l and b_k ,
3. Finding the best antecedent of b_k , denoted as $\zeta_{b_k-1}^*$, with the following equation:

$$\zeta_{b_k-1}^* = \operatorname{argmax}_{b_l} \left(u(Z_{\llbracket 1, b_l \rrbracket}^*) + u(\llbracket b_l, b_k - 1 \rrbracket) \right). \quad (5)$$

Finally, at the last iteration, the algorithm computes the best antecedent for $B + 1$, *i.e.* the last downbeat of the song. Then, recursively, the algorithm is able to backtrack the best antecedent of this antecedent, and so on and so forth back to the first bar of the song, thus providing the optimal segmentation. A graph visualization for a 4-bar example is presented in Figure 5. Pseudo-code for the CBM algorithm, assuming that the score function u is given, is detailed in the appendix (Algorithm 1).

In the end, for any bar b_k , the optimal segmentation up to b_k can be computed in $\mathcal{O}(b_k - 1)$ operations, *i.e.* parsing each antecedent only once. Hence, the solution algorithm boils down to $\mathcal{O}(\frac{B(B+1)}{2})$ evaluations, which corresponds to a polynomial complexity. In practice, we even limit the size of admissible segments to be at most 32 bars (set empirically), which further reduces the complexity.

3.2 SCORE FUNCTION

Finally, the segmentation problem boils down to the definition of the score function $u(\llbracket i, \zeta_{i+1} - 1 \rrbracket)$ for a segment. In the CBM algorithm and following (Sargent et al., 2016), the score of each segment is defined as a mixed score function, presented in Equation 6 as the balanced sum of two terms:

$$u(\llbracket \zeta_i, \zeta_{i+1} - 1 \rrbracket) = u^K(\llbracket \zeta_i, \zeta_{i+1} - 1 \rrbracket) - \lambda \rho(\zeta_{i+1} - \zeta_i). \quad (6)$$

The first term, $u^K(\llbracket \zeta_i, \zeta_{i+1} - 1 \rrbracket)$, is based on the homogeneity criterion, and is presented in Section 3.2.1; the second one, $\rho(\zeta_{i+1} - \zeta_i)$, is based on the regularity criterion, and is presented in Section 3.2.2. Parameter λ is a balancing parameter.

3.2.1 Block weighting kernels

The first term u^K of the score function in Equation 6 is obtained from the self-similarity values within a

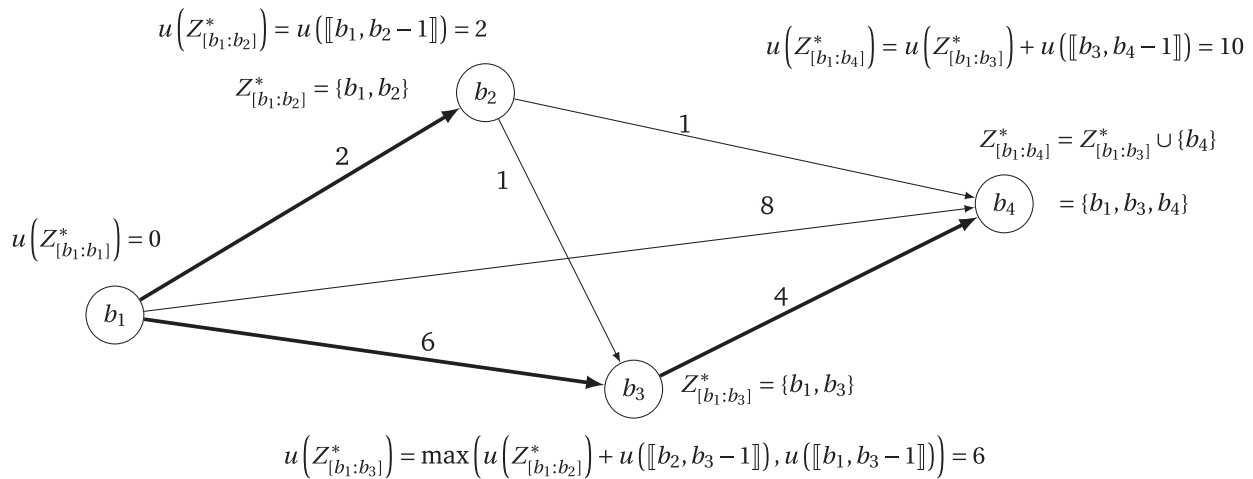


Figure 5 Example of computing an optimal segmentation with 4 bars.

segment. Practically, given a self-similarity matrix $A(X)$, the score $u^k(S_i)$ of segment $S_i = \llbracket \zeta_i, \zeta_{i+1} - 1 \rrbracket$ (of size $n = \zeta_{i+1} - \zeta_i$) is computed by evaluating the self-similarity values restricted to S_i , i.e. $A(X_{S_i}) = A(X_{[\zeta_i: \zeta_{i+1} - 1]})$. It can be understood as cropping the self-similarity $A(X)$ on this particular segment, around the diagonal.

The CBM algorithm aims at favoring the homogeneity of estimated segments, i.e. favoring sections composed of similar elements. Thus, the score function u^k is defined so as to measure the inner similarity of a segment. In practice, this is obtained through weighting local self-similarity values, by using a (fixed) weighting kernel matrix K , such as:

$$u^K : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$$

$$A(X_{S_i}) \mapsto \frac{1}{n} \sum_{k=1}^n \sum_{l=1}^n A(X_{S_i})_{kl} K_{kl}. \quad (7)$$

The kernel is called a “weighting kernel”. A first observation is that the weighting kernel needs to adapt to the size of the segment. A very simple kernel is a kernel matrix full of ones, i.e. $K = \mathbf{1}_{n \times n}$, resulting in a score function equal to the sum of every element in the self-similarity matrix, normalized by the size of the segment. The normalization by the size of the segment is meant to turn the squared dependence of the size of the segment in the number of self-similarity values (n^2) into a linear dependence. A linear dependence is desired as it ensures a length- n segment contributes similarly to the sum of segment scores as n segments of length 1.

The design of the weighting kernel defines how to transform bar similarities into segment homogeneity, which is of particular importance for segmentation. The remainder of this section presents two types of kernels, namely the “full” kernel and the “band” kernel. We consider that the main diagonal in the self-similarity matrix is not informative regarding the overall similarity in the segment, as its values are normalized to one. Hence, for every weighting kernel K used in the CBM algorithm, $K_{ii} = 0, \forall i$.

Full Kernel The first kernel is called the “full” kernel, because it corresponds to a kernel full of 1s (except on the diagonal where it is equal to 0). The full kernel captures the average value of similarities in this segment, excluding the self-similarity values. Practically, denoting as K^f the full kernel:

$$K_{ij}^f = \begin{cases} 1 & \text{if } i \neq j \\ 0 & \text{if } i = j \end{cases} \quad (8)$$

Hence, the score function associated with the full kernel is equal to:

$$u^{K^f}(S_i) = \frac{1}{n} \sum_{k=1}^n \sum_{l=1, l \neq k}^n A(X_{S_i})_{kl} K_{kl}^f = \frac{1}{n} \sum_{k=1}^n \sum_{l=1, l \neq k}^n A(X_{S_i})_{kl} \quad (9)$$

A full kernel of size 10 is presented in Figure 6.

Band Kernels A second class of kernels, called “band” kernels, are considered in order to emphasize short-term similarity. Indeed, in band kernels, the weighting score is computed on the pairwise similarities of a few bars in the segment only, depending on their temporal proximity: only close bars are considered. In practice, this can be obtained by defining a kernel with entries equal to 0, except on some upper- and sub-diagonals. The number of upper- and sub-diagonals is a parameter, corresponding to the maximal number of bars considered to evaluate the similarity, i.e. an upper bound on $|b_i - b_j|$ for a pair of bars (b_i, b_j) .

Hence, a band kernel is defined according to its number of bands, denoted as v , defining the v -band kernel K^{vb} such that:

$$K_{ij}^{vb} = \begin{cases} 1 & \text{if } 1 \leq |i - j| \leq v, \\ 0 & \text{otherwise } (i = j \text{ or } |i - j| > v). \end{cases} \quad (10)$$

Three band kernels, of size 10, are represented in Figure 7. Section 4 presents experiments which compare quantitatively the impact of the number of bands on the segmentation performance.

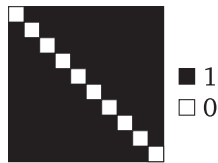


Figure 6 Full kernel of size 10.

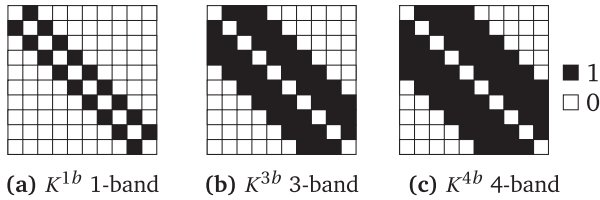


Figure 7 Band kernels, of size 10.

3.2.2 Penalty functions

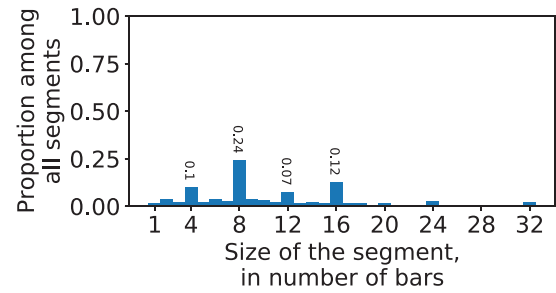
Sargent et al. (2016) extended the score function of Jensen (2006) to take into account both the homogeneity and the regularity criteria, resulting in Equation 6. In practice, this is obtained through defining a regularity penalty function $p(n)$, corresponding to the second term in Equation 6, and penalizing segments according to their size n , to favor particular sizes.

The penalty function is based on prior knowledge, and aims at enforcing particular sizes of segments, which are known to be typical in a number of music genres, notably Pop music. In particular, Figure 8 presents the distributions of the sizes of segments, in terms of number of bars, in the annotations of both RWC Pop and SALAMI datasets. It appears that some sizes of segments are much more frequent in the annotations. Hence, penalty functions p can be derived from these distributions.

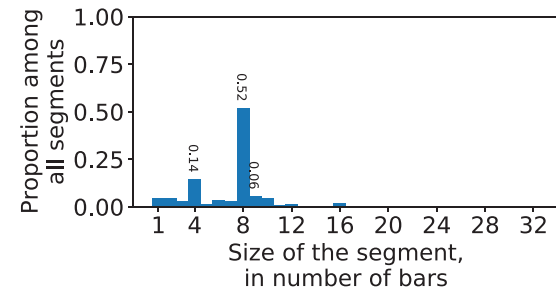
Two different penalty functions p are studied in this section, namely the “target-deviation” and “modulo” functions. In what follows, n denotes the size of the segment, i.e. $n = \zeta_{i+1} - \zeta_i$.

Target-Deviation Functions The first set of penalty functions, called “target-deviation” and denoted as p^{td} , is defined by Sargent et al. (2016). Target-deviation functions compute the difference between the size of the current estimated segment and a target size τ , raised to the power of a parameter α , i.e. $p^{td}(n) = |n - \tau|^\alpha$ where parameter α takes typical values in $\{0.5, 1, 2\}$. The target size is set by Sargent et al. (2016) to 32, to favor segments of size 32 beats, in line with their respective evaluations of most frequent segment sizes. In our barwise context,⁶ $\tau = 8$, which is the most frequent segment size in both RWC Pop and SALAMI datasets.

This penalty function is adapted to enforce one size in particular, and tends to disadvantage all the others. Hence, this function is adapted to datasets where one size is predominant, which seems true for RWC Pop with MIREX 10 annotations (more than half of the segments in the annotation are of size 8 bars), but not so definite for the SALAMI dataset, where the segment sizes are



(a) SALAMI-test.



(b) RWC Pop.

Figure 8 Distribution of segment sizes in terms of number of bars, in the annotations.

more balanced between 4, 8, 12 and 16, as presented in Figure 8. In particular, segments of size 16 are strongly penalized ($|8-16|^\alpha = 8^\alpha$).

Modulo function The second set of penalty functions, called “modulo functions”, is designed to favor particular segment sizes, directly based on prior knowledge. In this study, we only present the “modulo 8” function $p^{m8}(n)$ based on both RWC Pop and SALAMI annotations. Indeed, in both datasets, most segments are of size 8, and the remaining segments are generally of size 4, 12 or 16. Finally, outside of these sizes, even segments are more frequent than segments of odd sizes. Hence, the modulo 8 function models this distribution, as:

$$p^{m8}(n) = \begin{cases} 0 & \text{if } n=8 \\ \frac{1}{4} & \text{else, if } n \equiv 0 \pmod{4} \\ \frac{1}{2} & \text{else, if } n \equiv 0 \pmod{2} \\ 1 & \text{otherwise} \end{cases} \quad (11)$$

Penalty values for the different cases were set quite intuitively, and would benefit from further investigation.

In order to mitigate both the weighting score function u^k and the penalty function p , we implemented an additional normalization step based on the weighted values obtained in each song, resulting in the score function defined in Equation 12.

$$u(\llbracket \zeta_i, \zeta_{i+1} - 1 \rrbracket) = u^k(\llbracket \zeta_i, \zeta_{i+1} - 1 \rrbracket) - u_{\max}^{k8} \lambda p(\zeta_{i+1} - \zeta_i), \quad (12)$$

In Equation 12, u_{\max}^{k8} is the maximal weighting value obtained by sliding a kernel of size 8 on this self-similarity matrix, i.e. the highest score among all possible segments

of size 8. This size of 8 for the kernel is chosen as the most frequent segment size in terms of number of bars in both RWC Pop and SALAMI datasets, as presented in Figure 8. Parameter λ is a constant parameter, which is fitted as detailed in Section 4.

Finally, in the CBM algorithm, the score of each segment is defined as in Equation 12. The first term, $u^k(\square_{\zeta_i, \zeta_{i+1}} - 1)$, is a weighting score, measuring the self similarity of the segment. The second term, $p(\zeta_{i+1} - \zeta_i)$, penalizes or favors the segment depending on its size. Both these scores are subject to design choices, which are studied and compared in the subsequent section.

4. EXPERIMENTS

4.1 EVALUATION METRICS

The quality of the estimation obtained with the CBM algorithm is evaluated with the Hit-Rate metrics, comparing a set of estimated boundaries with a set of annotations by intersecting them with respect to a tolerance t (Ong and Herrera, 2005; Turnbull et al., 2007). In practice, given two sets of boundaries Z^e and Z^a (respectively the sets of estimated and annotated boundaries), an estimated boundary $\zeta_i^e \in Z^e$ is considered correct if it is close enough to an annotated boundary $\zeta_j^a \in Z^a$ (“close enough” meaning that the gap is no larger than the tolerance t), i.e. if $\exists \zeta_j^a \in Z^a$ such that $|\zeta_i^e - \zeta_j^a| \leq t$. Each estimated boundary can be coupled with a maximum of one annotated boundary, and *vice versa*. The set of correct boundaries subject to the tolerance t , denoted as C_t , contains at most as many elements as the annotations or the estimations, i.e. $0 \leq |C_t| \leq \min(|Z^e|, |Z^a|)$. In case of perfect concordance between Z^e and Z^a , $C_t = Z^e = Z^a$. In practice, the concordance of C_t with Z^e and Z^a is evaluated by the precision P_t , recall R_t and F-measure F_t :

- $P_t = \frac{|C_t|}{|Z^e|}$, i.e. the proportion of accurately estimated boundaries among the total number of estimated boundaries.
- $R_t = \frac{|C_t|}{|Z^a|}$, i.e. the proportion of accurately estimated boundaries among the total number of annotated boundaries.
- $F_t = \frac{2P_tR_t}{P_t+R_t}$ is the harmonic mean of both aforementioned measures. The harmonic mean is less sensitive to large values than the arithmetic (standard) mean, and is conversely more strongly penalized by low values. Hence, a high F-measure requires both a high recall and a high precision.

These metrics are computed using the *mir_eval* toolbox (Raffel et al., 2014).

4.1.1 Tolerances in absolute time

In the boundary retrieval subtask, conventions for the tolerance values are 0.5s (Turnbull et al., 2007) and 3s

(Ong and Herrera, 2005). The 3-second tolerance, citing Ong and Herrera (2005), is justified as being equal to “approximately 1 bar for a song of quadruple meter [NB: 4 beats per bar, e.g. $\frac{4}{4}$ metric] with 80 bpm in tempo”, while the 0.5 second tolerance is within the order of magnitude corresponding to the beat. In this work, we use both tolerance values to compare our algorithm with the standard algorithms, leading to 6 metrics $P_{0.5s}$, $R_{0.5s}$, $F_{0.5s}$ and P_{3s} , R_{3s} , F_{3s} .

4.1.2 Barwise-aligned tolerances

In this work, estimated boundaries are located on downbeat estimates, as explained and motivated in Section 2. In that sense, rather than evaluating the estimates in absolute time, we align each annotation with the closest estimated downbeat, leading to barwise-aligned annotations. This allows us to introduce additional metrics: P_{0bar} , R_{0bar} , F_{0bar} and P_{1bar} , R_{1bar} , F_{1bar} . The first three metrics (e.g. F_{0bar}) consider that the tolerance is set to 0 bars, i.e. expecting estimates and annotations to fall precisely on the same downbeat, and the latter three metrics (e.g. F_{1bar}) set the tolerance to exactly one bar between estimates and annotations. In particular, these metrics will be used to compare different settings of our algorithm.

4.2 PARAMETRIZATION OF THE ALGORITHM

The CBM algorithm is evaluated on the boundary retrieval task on the entire RWC Pop dataset (Goto et al., 2002), and on the test subset of SALAMI (Smith et al., 2011), defined by Ullrich et al. (2014). The three similarity functions defined in Section 2.2 are used to compute the self-similarity matrices, namely Cosine, Autocorrelation and RBF. The CBM algorithm itself is subject to the choice of kernel, and particularly to the number of bands when using a band kernel. In addition, the score function depends on the design of the penalty function. Rather than studying all of these parameters at the same time, experiments focus on each aspect independently. In particular, the experiments aim at answering the following three questions:

- Which similarity function is the most suitable for boundary retrieval in our context?
- Which weighting kernel is the most suitable for boundary retrieval in our context?
- Which penalty function is the most suitable for boundary retrieval in our context?

Each question is addressed sequentially, and the conclusion of each question serves as the basis to study the next ones.

4.2.1 Train/test datasets

These questions are addressed by comparing several parameters in a train/test fashion: a subset of the SALAMI dataset, called “SALAMI-train”, is used to

evaluate several parameters, and the best one in this subset is evaluated on the remainder of the SALAMI dataset, called “SALAMI-test”, and on the entire RWC Pop dataset. The division between SALAMI-train and SALAMI-test is defined by Ullrich et al. (2014), based on the MIREX evaluation dataset. The details are available online,⁷ and are uploaded along with experimental Notebooks on the open-source dataset.⁸ The SALAMI-train dataset contains 849 songs, and the SALAMI-test dataset contains 485 songs.⁹ The entire RWC Pop dataset contains 100 songs, resulting in a total of 585 songs for testing.

4.2.2 Self-similarity matrices

Firstly, we study the impact of the design of the similarity function on the performance of the CBM algorithm. To do so, we use the CBM algorithm with the full kernel, as it does not need the fitting of the number of bands, and we do not use a penalty function. The boundary retrieval performance is presented in Table 4 for the train dataset.

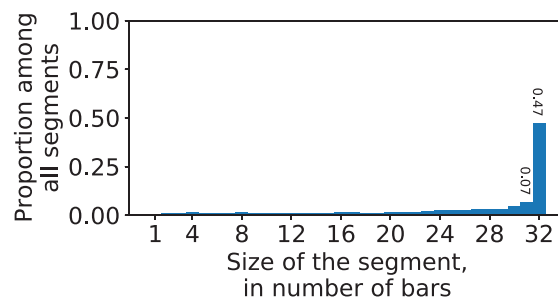
The RBF self-similarity is the best-performing self-similarity in terms of F-measure (with both tolerances), hence suggesting a better boundary estimation on average than the other similarity functions. The results obtained with the RBF similarity function on the test datasets are presented in Table 5.

The precision/recall trade-offs depend on the self-similarity matrices, and deserve to be studied to give further information on the quality of the estimated segmentations. The Cosine self-similarity exhibits a higher precision than recall on average, which suggests an under-segmentation, *i.e.* estimating too few boundaries. Conversely, the Autocorrelation self-similarity results in a higher recall than precision, suggesting over-segmentation. The RBF self-similarity performance is more balanced between both metrics.

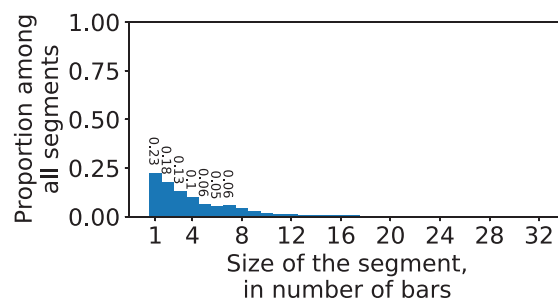
These conclusions can be confirmed by studying the distribution of the sizes of the estimated segments, as presented in Figure 9 on the SALAMI-train dataset. These distributions must be compared with the distribution of segment sizes in the annotations, presented in Figure 8.

The distribution of segment sizes with the RBF self-similarity is visually the closest one to the distribution of

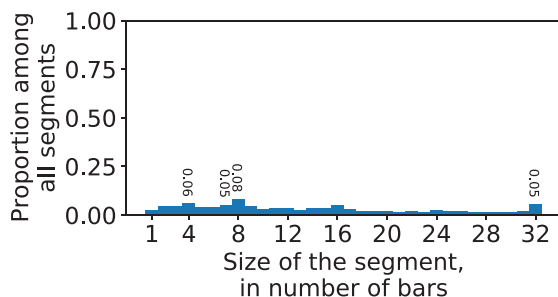
annotations, which we confirm numerically by studying the Kullback-Leibler (KL) divergences between the distribution of the sizes of the estimated segments and of the annotated ones. The KL-divergences are respectively equal to 2.25, 0.85 and 0.35 for the Cosine, Autocorrelation and RBF similarity functions. Again, this suggests that the RBF similarity function is the most suitable.



(a) Cosine self-similarity.



(b) Autocorrelation self-similarity.



(c) RBF self-similarity.

Figure 9 Distribution of segment sizes, with the full kernel, according to the self-similarity matrix. Results on the SALAMI-train dataset.

Self-similarity	P_{0bar}	R_{0bar}	F_{0bar}	P_{1bar}	R_{1bar}	F_{1bar}
Cosine	50.83%	30.82%	36.77%	62.80%	37.72%	45.19%
Autocorrelation	32.59%	64.69%	41.30%	42.10%	83.73%	53.41%
RBF	50.27%	45.38%	45.84%	64.79%	58.81%	59.30%

Table 4 Boundary retrieval performance with the different self-similarities on the train dataset (Full kernel, no penalty function).

Dataset	P_{0bar}	R_{0bar}	F_{0bar}	P_{1bar}	R_{1bar}	F_{1bar}
SALAMI - test	48.52%	48.65%	46.68%	62.76%	63.09%	60.51%
RWC Pop	60.72%	53.61%	56.01%	77.68%	67.62%	71.09%

Table 5 Boundary retrieval performance with the RBF self-similarity, on both test datasets (Full kernel, no penalty function).

4.2.3 Block weighting kernels

Secondly, an important parameter in the CBM algorithm is the design of the kernel. We thus compare the full kernel with band kernels, the number of bands varying from 1 to 16. Results on the SALAMI-train dataset, computed on the RBF self-similarity matrices, and focusing on the F-measures, are presented in Figure 10. The 7-band kernel stands out as the best-performing kernel, even if performance is close to the 15-band kernel.

The differences in performance between the different kernels may be explained by Figure 11, which presents the distribution of segment sizes according to the number of bands. The 7-band kernel leads to a majority of estimated segments of size 8 (more than 50%, twice as much as in the annotations), which is the most common segment size in the annotation, while the 15-band kernel mostly computes segments of size 16, and the full kernel is well distributed across the different segment sizes. The

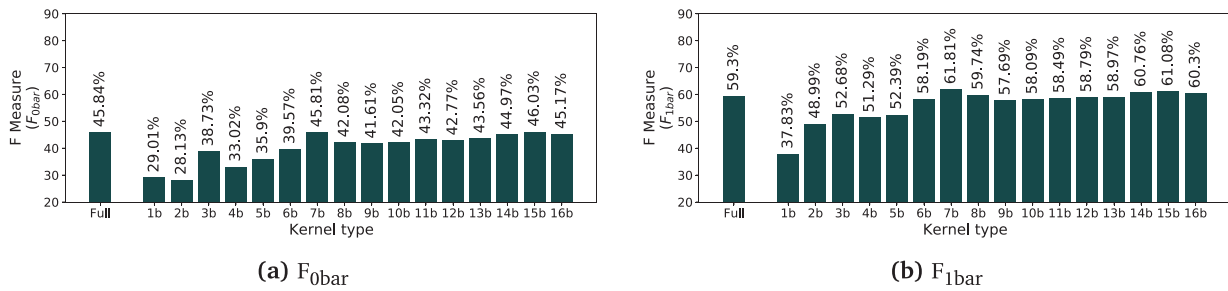


Figure 10 Boundary retrieval performance (F-measures only) according to the full and band kernels (with different numbers of bands). Results on the train dataset with RBF self-similarity matrices.

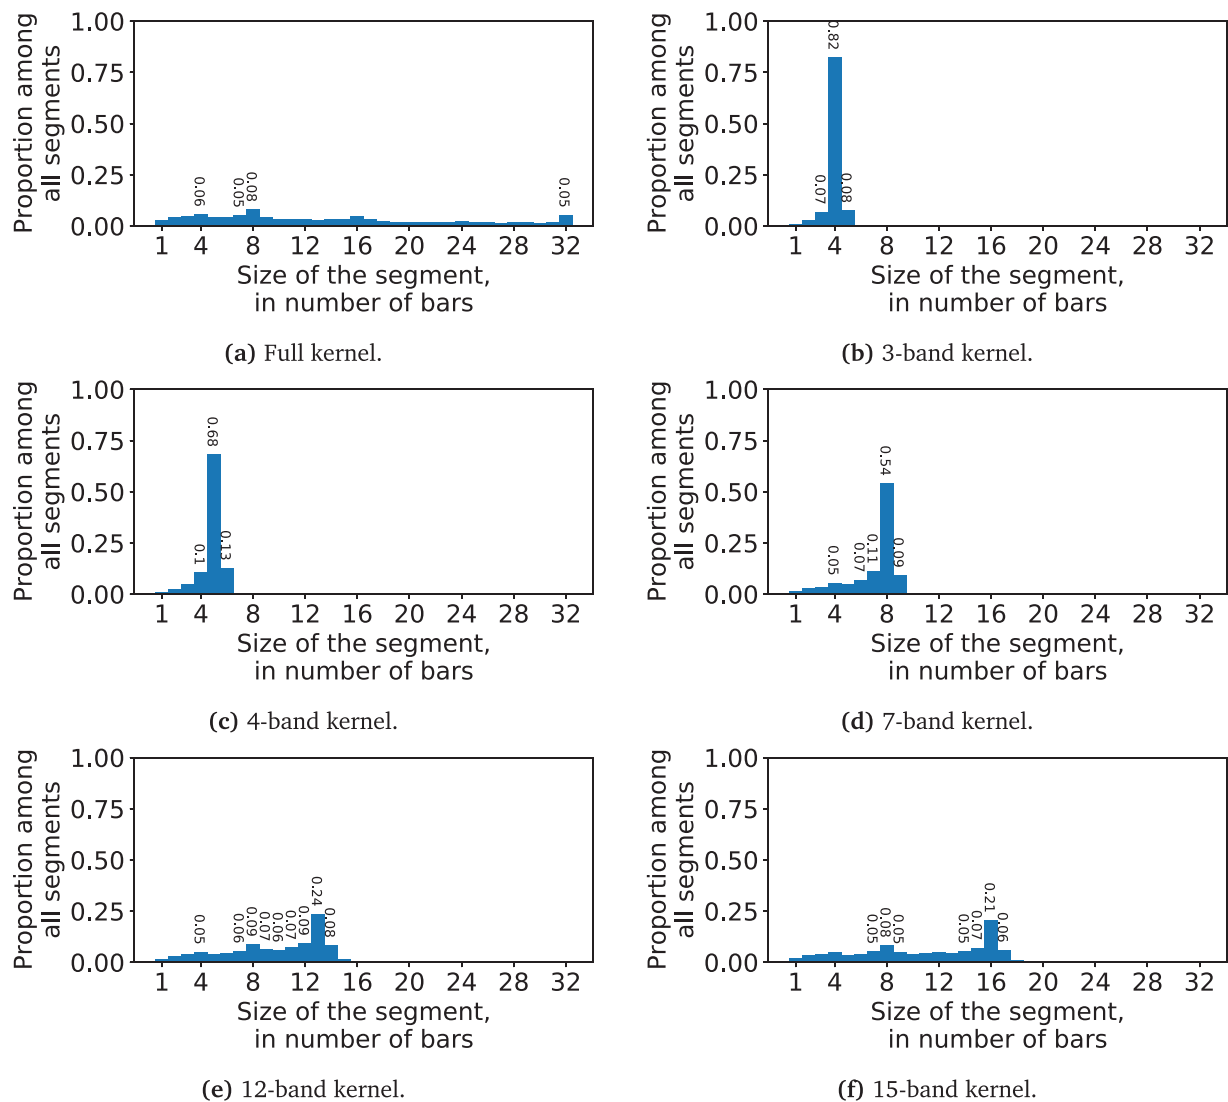


Figure 11 Distribution of estimated segment sizes, according to different kernels, on the train dataset.

annotations are mostly composed of segments of size 8, then 4, 12 and 16. Hence, while the 7-band kernel does not accurately represent the annotations, it obtains better boundary retrieval performance than the other ones, indicating that this latter distribution is beneficial to the boundary estimation overall.

As an additional conclusion, the number of bands in the kernel largely influences the distribution of segment sizes, in particular the most frequent segment size. As a general trend, it seems that a kernel with v bands favors segments of size $v + 1$. We assume that this behavior stems from the fact that, for a v -band kernel and a large segment of size $n > v$, the number of elements equal to 0 is large, but the normalization remains adapted to kernels with n^2 values. We found in practice that this effect could be dampened by normalizing the score associated with each kernel by the number of nonzero values plus the number of elements in the diagonal instead of the size of the kernel, as (Shiu et al., 2006) did, but this resulted in all kernels performing similarly to the full kernel, hence performing worse than the 7-band one.

Finally, as the 7-band kernel is the best-performing one, we fixed this kernel for both test datasets. Results obtained with this kernel are presented in Table 6.

4.2.4 Penalty functions

Finally, the last experiments focus on the penalty functions. In this set of experiments, we compare the

target deviation functions, with $\alpha \in \{0.5, 1, 2\}$, with the modulo 8 function. The CBM algorithm is parametrized with the 7-band kernel, and is applied on the RBF self-similarity matrices. The parameter λ , balancing the penalty function, takes values between $\frac{1}{100}$ and $\frac{2}{10}$, with a step of $\frac{1}{100}$. This parameter is fitted on the SALAMI-train dataset. Results are presented in Table 7.

The modulo 8 function appears to slightly improve boundary retrieval performance for the metrics with a tolerance of 0 bar, indicating a more accurate estimation, but results with a tolerance of 1 bar are not strongly impacted by the choice of the penalty function. Results are close between the different penalty functions, except for the target deviation with a large α , which results in worse performance than the other conditions.

Overall, it seems that the modulo 8 function is the most suitable penalty function to estimate segments accurately. Hence, we use this penalty function for the test results, presented in Table 8, with parameter $\lambda = 0.04$, as optimized on the SALAMI-train dataset.

4.2.5 Experimental conclusions

In light of these results, we finally sum up the situation regarding the choice of settings for the CBM algorithm.

1. *In our context, the RBF self-similarity matrix is the most suitable self-similarity matrix for boundary retrieval.*

Dataset	P _{0bar}	R _{0bar}	F _{0bar}	P _{1bar}	R _{1bar}	F _{1bar}
SALAMI - test	37.24%	59.80%	44.33%	50.38%	80.52%	59.88%
RWC Pop	59.41%	68.19%	62.82%	75.53%	86.56%	79.81%

Table 6 Boundary retrieval performance with the 7-band kernel, on both test datasets (RBF self-similarity, no penalty function).

Penalty function	Best λ	P _{0bar}	R _{0bar}	F _{0bar}	P _{1bar}	R _{1bar}	F _{1bar}
Without penalty	-	40.26%	57.38%	45.81%	54.26%	77.67%	61.81%
Target deviation	$\alpha = \frac{1}{2}$	0.01	40.38%	57.36%	45.88%	54.37%	61.84%
	$\alpha = 1$	0.01	40.45%	56.98%	45.81%	54.61%	61.89%
	$\alpha = 2$	0.01	39.75%	54.32%	44.43%	54.93%	61.46%
Modulo 8	0.04	41.04%	58.34%	46.63%	54.25%	77.44%	61.72%

Table 7 Boundary retrieval performance depending on the penalty function, for the SALAMI-train dataset, with the RBF self-similarity and the 7-band kernel.

Dataset	P _{0bar}	R _{0bar}	F _{0bar}	P _{1bar}	R _{1bar}	F _{1bar}
SALAMI - test	38.36%	60.96%	45.44%	50.76%	80.51%	60.09%
RWC Pop	62.11%	70.05%	65.17%	77.35%	86.95%	81.02%

Table 8 Boundary retrieval performance with the modulo 8 penalty function ($\lambda = 0.04$), on both test datasets (RBF self-similarity, 7-band kernel).

2. In our context, the 7-band kernel is the most suitable kernel for boundary retrieval.
3. In our context, the modulo 8 penalty function is the most suitable penalty function for boundary retrieval.

4.2.6 Metrics with tolerance in absolute time

As mentioned in Section 4.1, standard metrics for boundary retrieval performance consider the tolerance in absolute time (e.g. $F_{0.5s}$ and F_{3s} metrics), while we opted for boundary-aligned metrics in our experiments. Hence, Table 9 compares the boundary retrieval performance obtained when the tolerance is defined relatively to the bars and in absolute time, which allows to compare with state-of-the-art algorithms. Boundary retrieval performance is almost equivalent on the RWC Pop dataset, and slightly altered for the metrics with short tolerances on the SALAMI dataset (F_{0bar} and $F_{0.5s}$). These discrepancies may be explained by the less precise downbeat alignment of annotations in the SALAMI dataset, presented in Table 1. Overall though, results remain similar, which tends to confirm the hypothesis of Ong and Herrera (2005) that a tolerance of 3 seconds corresponds approximately to a tolerance of 1 bar.

Dataset	F_{0bar}	$F_{0.5s}$	F_{1bar}	F_{3s}
SALAMI-test	45.44%	42.00%	60.09%	60.61%
RWC Pop	65.17%	64.44%	81.02%	80.64%

Table 9 Boundary retrieval performance, comparing the F-measures with tolerance expressed barwise and in absolute time.

4.3 COMPARISON WITH STATE-OF-THE-ART ALGORITHMS

We compare the boundary retrieval performance obtained by the Foote-TF (introduced in Section 2.3.3) and the CBM algorithms with state-of-the-art algorithms. The performance of the CBM algorithm is obtained using the hyperparameters learned in Section 4.2.

This work considers seven different algorithms as state-of-the-art, categorized as either unsupervised or supervised algorithms, i.e. algorithms that either estimate boundaries without the use of training examples or analyze annotated examples before making predictions: four unsupervised algorithms (Foote, 2000; McFee and Ellis, 2014a; Serrà et al., 2014; McCallum, 2019) and three supervised algorithms (Grill and Schlüter, 2015; Wang et al., 2021; Salamon et al., 2021). We additionally use previous work on the CBM algorithm (Marmoret et al., 2022b) as a baseline.

All state-of-the-art algorithms use beat-aligned features, except Grill and Schlüter (2015), who use a fixed hop length and Wang et al. (2021), who use downbeat-aligned features. Results for Foote (2000); McFee and Ellis, (2014a); Serrà et al. (2014) are computed with the MSAF toolbox (Nieto and Bello, 2016), and realigned on downbeats in post-processing. Results for the CNN (Grill and Schlüter, 2015) are extracted from the 2015 MIREX contest. Results for McCallum (2019); Wang et al. (2021); Salamon et al. (2021) are copied from the respective articles.

Figures 12 and 13 compare the results obtained with the CBM and the Foote-TF algorithms with

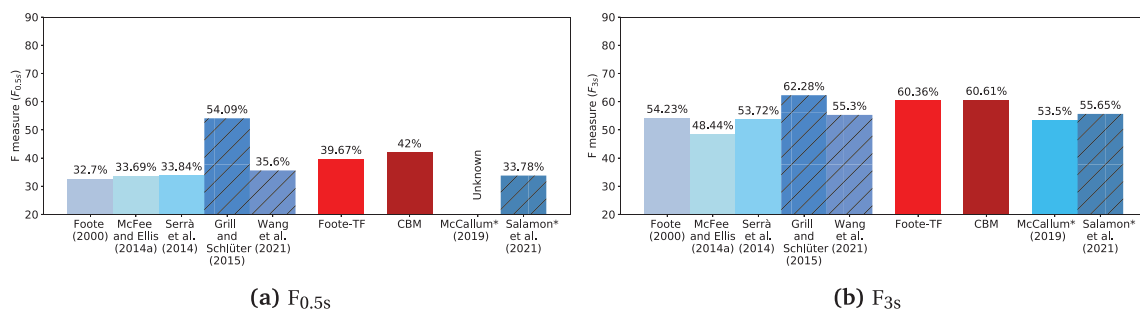


Figure 12 Boundary retrieval performance of the CBM algorithm on the SALAMI dataset, compared to state-of-the-art algorithms. Hatched bars correspond to supervised algorithms. The star * represents algorithms where the evaluation subset is not exactly the same as ours, thus preventing accurate comparison.

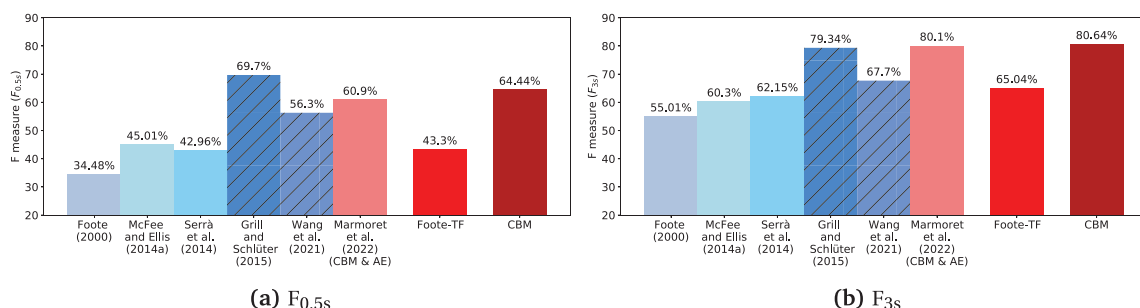


Figure 13 Boundary retrieval performance of the CBM algorithm on the RWC Pop dataset, compared to state-of-the-art algorithms. Hatched bars correspond to supervised algorithms.

those of the state-of-the-art algorithms.¹⁰ In this comparison, the CBM algorithm globally outperforms the other unsupervised segmentation methods, most of the supervised algorithms, and is competitive for the metric F_{3s} with the global (supervised) state-of-the-art (Grill and Schlüter, 2015). These results are promising and show the potential of the CBM algorithm, which is performing well despite its relative simplicity. Additionally, results of the CBM algorithm are on par with those of Foote-TF on the SALAMI dataset, showcasing interest for the Barwise TF representation.¹¹

4.4 CBM: BAR-SCALE VS. BEAT-SCALE

In order to distinguish the impact of barwise-alignment and of the CBM algorithm itself on the segmentation results, we compare results obtained on the Barwise TF matrix, presented above, with results obtained on the “Beatwise TF matrix”, *i.e.* the equivalent of the Barwise TF matrix at the beat-scale. Following the intuition that most bars in Western modern music are composed of 4 beats per bar, the Beatwise TF matrix is sampled with $T = 24$ samples per beat, *i.e.* $\frac{96}{4}$. Beats are estimated with the algorithm of Böck et al. (2019), implemented in the *madmom* toolbox, as in Section 2.3. For fairer comparison, the results at both beatwise and barwise scales are computed without penalty function. Corresponding results are presented in Tables 10 and 11.

Results confirm that barwise-alignment is beneficial for the performance of the CBM algorithm, especially on the RWC Pop dataset, but the CBM algorithm with beat-aligned features obtains similar or better performance than the unsupervised state-of-the-art algorithms (Foote, 2000; McFee and Ellis, 2014a; Serrà et al., 2014; McCallum, 2019), and is competitive with the supervised algorithm of Wang et al. (2021) on both datasets for the metric F_{3s} .

5. CONCLUSIONS

This article presented the CBM algorithm for performing Music Structure Analysis on audio signals, where boundaries between musical sections are computed by maximizing the homogeneity of each segment composing the segmentation, using dynamic programming under a penalty function. Moreover, barwise processing of music is shown to increase segmentation performance, using the Barwise TF matrix. This work has also investigated several metrics to represent similarities between pairs of bars in a song. While the CBM algorithm has room for improvement, it achieves a level of performance which is competitive to the state-of-the-art, and therefore appears as a meaningful approach to investigate a variety of music representations without needing large collections of training data.

The design of the kernel clearly impacts boundary retrieval performance. Hence, future work could focus on studying alternative types of kernels. The kernel values could depend on the particular song or dataset considered, or follow particular statistical distributions. Of particular interest could be the learning of such kernels instead of an (empirical) definition. These latter comments are also valid for the penalty functions, whose values were set quite empirically, and which would benefit from deeper investigation. The number of bands in the weighting kernels seems to enforce particular segment sizes. This effect can be mitigated with normalization, or, conversely, further exploited, for instance by using different kernels concurrently, each one accounting for a different level of structure, hence studying segmentation hierarchically.

Weighting kernels presented in this article focus on the homogeneity of each segment, but other kernels could be considered in order to account for repetition in the song. In fact, the proposed framework is highly customizable with respect to weighting kernels, and could be adapted

SALAMI	$P_{0.5s}$	$R_{0.5s}$	$F_{0.5s}$	P_{3s}	R_{3s}	F_{3s}
Beatwise (cosine, 63-band kernel)	35.90%	41.61%	37.36%	55.75%	64.52%	58.03%
Barwise (RBF, 7-band kernel)	34.49%	54.56%	41.04%	50.70%	80.78%	60.51%

Table 10 CBM algorithm, performed on Barwise TF matrix vs. Beatwise TF matrix, on the SALAMI-test dataset. For fairer comparison, results at both scales are computed without penalty function.

RWC Pop	$P_{0.5s}$	$R_{0.5s}$	$F_{0.5s}$	P_{3s}	R_{3s}	F_{3s}
Beatwise (cosine, 63-band kernel)	46.22%	44.38%	44.57%	72.54%	68.85%	69.51%
Barwise (RBF, 7-band kernel)	59.09%	67.13%	62.28%	75.17%	85.90%	79.47%

Table 11 CBM algorithm, performed on Barwise TF matrix vs. Beatwise TF matrix, on RWC Pop. For fairer comparison, results at both scales are computed without penalty function.

to the expected shape of segments. In particular, we expect that bridging this work with previous work (Foote, 2000) could further enhance performance.

6. REPRODUCIBILITY

All the code used in this article is contained in the open-source toolbox (Marmoret et al., 2022a), along with experimental Notebooks used to compute the experimental results.⁸

A ALGORITHM, IN DETAILS

The detailed algorithm, in pseudo-code, is presented hereafter.

Algorithm 1 CBM algorithm, computing the optimal segmentation given a score function $u()$.

Input: Bars $\{b_k \in \llbracket 1, B \rrbracket\}$, score function u
Output: Optimal segmentation $Z^* = \{\zeta\}$

```

 $Z^* = \{1, B + 1\}$ 
 $A^* = []$ 
/* Array storing the optimal antecedents for
   every bar (empty at initialization). */

 $\mathcal{U}^* = [0]$ 
/* Array storing the optimal segmentation up
   to each bar (set to  $\mathcal{U}^*[1] = 0$  at
   initialization). */
for  $b_k = 2, \dots, B + 1$  do
   $\zeta = b_k$ 
  /* Considering bar  $b_k$  as current boundary.
   */

   $\zeta_{-1}^* = \arg \max_{1 \leq \zeta_j < \zeta} (\mathcal{U}^*[\zeta_j] + u(\llbracket \zeta_j, \zeta - 1 \rrbracket))$ 
  /* Finding the best antecedent for the
   current boundary  $\zeta$  with Equation 5. */

   $A^*[\zeta] = \zeta_{-1}^*$ 
  /* Storing the best antecedent for  $\zeta$ . */

   $\mathcal{U}^*[\zeta] = \mathcal{U}^*[\zeta_{-1}^*] + u(\llbracket \zeta_{-1}^*, \zeta - 1 \rrbracket)$ 
  /* Computing and storing the optimal
   segmentation score up to  $\zeta$ . */
end
 $\zeta = B + 1$ 
while  $A^*[\zeta] \neq 1$  do
  /* Recursively backtracking all optimal
   antecedents, from the last to the first
   bar. */

   $\zeta_{-1}^* = A^*[\zeta]$ 
   $Z^* = Z^* \cup \{\zeta_{-1}^*\}$ 
  /* Searching for the best antecedent and
   adding it to the optimal segmentation.
   */

   $\zeta = \zeta_{-1}^*$ 
  /* Iterating the process with the current
   best antecedent. */
end

```

Practically, the advantage of the algorithm is to be able to store in memory both the optimal antecedent for each bar and the scores of the optimal segmentation up to each bar when they are computed for the first time, respectively denoted as the arrays A^* and \mathcal{U} in Algorithm 1.

NOTES

- 1 www.music-ir.org/mirex/wiki/2016:Audio_Downbeat_Estimation_Results.
- 2 In this experiment, the Barwise TF matrix is computed on the same features as in the implementation of the algorithm of Foote (2000) in MSAF (i.e. chroma features).
- 3 As the song contains B bars, $B + 1$ represents the end of the last bar, i.e. the last downbeat of the song.
- 4 As each set of boundaries must contain the first and last downbeats of the song, at most 2^{B-1} sets of boundaries can be obtained.
- 5 In details, both Jensen (2006) and Sargent et al. (2016) introduced the optimal segmentation as the minimum of a cost function, when it is rather defined here as a maximum. It actually depends on the way of conceiving the score function u , and, in particular, by defining a cost function equal to the inverse of the score function u , both problems are equivalent.
- 6 Note that 8 bars containing 4 beats each leads to 32 beats. 4 beats per bar is a frequent value for Western Pop music.
- 7 jan-schlueter.de/pubs/2014_ismir/.
- 8 https://gitlab.imt-atlantique.fr/a23marmo/autosimilarity_segmentation/-/tree/TISMIR.
- 9 There is a slight difference in the number of songs in the dataset due to some songs being missing in our version of the SALAMI dataset.
- 10 A careful reader will notice that results of Foote (2000) in Figures 12 and 13 are not exactly the same as those obtained in Tables 2 and 3. We explain these discrepancies by the use of different beatwise estimation algorithms: the *madmom* toolbox for Figures 12 and 13, and the original implementation in the MSAF toolbox for Tables 2 and 3.
- 11 We recall here that the Barwise TF matrix of the Foote-TF algorithm uses the parameters of Foote (2000), i.e. Chromagrams and Cosine self-similarity matrix. Still, Foote's algorithm implementation in Nieto and Bello (2016) also uses pre-processing steps which benefit the algorithm, and could benefit the CBM algorithm too. Studying these types of bridges between both algorithms is left to future work.

FUNDING INFORMATION

This work is partly supported by ANR JCJC project LoRAiA (ANR-20-CE23-0010).

COMPETING INTERESTS

AM has worked with both Simon Leglaive and Nancy Bertin as part of this work.

JC is partly funded by ANR JCJC project LoRAiA ANR-20-CE23-0010.

FB is a member of the ISCA Advisory Council (IAC), and Editor-in-Chief of the Journal "Speech Communication".

AUTHOR AFFILIATIONS

Axel Marmoret  orcid.org/0000-0001-6928-7490

Univ. Rennes 1, Inria, CNRS, IRISA, France;
IMT Atlantique, Lab-STICC, Brest, France

Jérémy E. Cohen  orcid.org/0000-0001-8319-8566

CREATIS, Univ Lyon, CNRS, France

Frédéric Bimbot

Univ. Rennes 1, Inria, CNRS, IRISA, France

REFERENCES

- Bellman, R.** (1952). On the theory of dynamic programming. *Proceedings of the National Academy of Sciences*, 38(8):716–719. DOI: <https://doi.org/10.1073/pnas.38.8.716>
- Bengio, Y., Courville, A., and Vincent, P.** (2013). Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828. DOI: <https://doi.org/10.1109/TPAMI.2013.50>
- Böck, S. and Davies, M. E.** (2020). Deconstruct, analyse, reconstruct: How to improve tempo, beat, and downbeat estimation. In *International Society for Music Information Retrieval Conference (ISMIR)*, pages 574–582.
- Böck, S., Davies, M. E., and Knees, P.** (2019). Multitask learning of tempo and beat: Learning one to improve the other. In *International Society for Music Information Retrieval Conference (ISMIR)*, pages 486–493.
- Böck, S., Korzeniowski, F., Schlüter, J., Krebs, F., and Widmer, G.** (2016a). Madmom: A new Python audio and music signal processing library. In *Proceedings of the 24th ACM International Conference on Multimedia*, pages 1174–1178. DOI: <https://doi.org/10.1145/2964284.2973795>
- Böck, S., Krebs, F., and Widmer, G.** (2016b). Joint beat and downbeat tracking with recurrent neural networks. In *International Society for Music Information Retrieval Conference (ISMIR)*, pages 255–261.
- Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C.** (2009). *Introduction to Algorithms*. MIT press, 3rd edition.
- de Berardinis, J., Vamvakaris, M., Cangelosi, A., and Coutinho, E.** (2020). Unveiling the hierarchical structure of music by multi-resolution community detection. *Transactions of the International Society for Music Information Retrieval*, 3(1):82–97. DOI: <https://doi.org/10.5334/tismir.41>
- Foote, J.** (2000). Automatic audio segmentation using a measure of audio novelty. In *IEEE International Conference on Multimedia and Expo*, pages 452–455. DOI: <https://doi.org/10.1109/ICME.2000.869637>
- Fuentes, M., McFee, B., Crayencour, H. C., Essid, S., and Bello, J. P.** (2019). A music structure informed downbeat tracking system using skip-chain conditional random fields and deep learning. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 481–485. DOI: <https://doi.org/10.1109/ICASSP.2019.8682870>
- Goto, M., Hashiguchi, H., Nishimura, T., and Oka, R.** (2002). RWC Music Database: Popular, Classical and Jazz Music Databases. In *International Conference on Music Information Retrieval (ISMIR)*, pages 287–288.
- Grill, T. and Schlüter, J.** (2015). Music boundary detection using neural networks on combined features and two-level annotations. In *International Society for Music Information Retrieval Conference (ISMIR)*, pages 531–537. DOI: <https://doi.org/10.1109/EUSIPCO.2015.7362593>
- Hung, Y.-N., Wang, J.-C., Song, X., Lu, W.-T., and Won, M.** (2022). Modeling beats and downbeats with a time-frequency transformer. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 401–405. DOI: <https://doi.org/10.1109/ICASSP43922.2022.9747048>
- Jensen, K.** (2006). Multiple scale music segmentation using rhythm, timbre, and harmony. *EURASIP Journal on Advances in Signal Processing*, 2007:1–11. DOI: <https://doi.org/10.1155/2007/73205>
- Marmoret, A., Cohen, J., and Bimbot, F.** (2022a). as_seg: Module for computing and segmenting autosimilarity matrices. https://gitlab.inria.fr/amarmoret/autosimilarity_segmentation.
- Marmoret, A., Cohen, J. E., Bertin, N., and Bimbot, F.** (2020). Uncovering audio patterns in music with nonnegative Tucker decomposition for structural segmentation. In *International Society for Music Information Retrieval Conference (ISMIR)*, pages 788–794.
- Marmoret, A., Cohen, J. E., and Bimbot, F.** (2022b). Barwise compression schemes for audio-based music structure analysis. In *19th Sound and Music Computing Conference (SMC 2022)*.
- Mauch, M., Noland, K. C., and Dixon, S.** (2009). Using musical structure to enhance automatic chord transcription. In *International Society for Music Information Retrieval Conference (ISMIR)*, pages 231–236.
- McCallum, M. C.** (2019). Unsupervised learning of deep features for music segmentation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 346–350. DOI: <https://doi.org/10.1109/ICASSP.2019.8683407>
- McFee, B. and Ellis, D.** (2014a). Analyzing song structure with spectral clustering. In *International Society for Music Information Retrieval Conference (ISMIR)*, pages 405–410.
- McFee, B. and Ellis, D. P.** (2014b). Learning to segment songs with ordinal linear discriminant analysis. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5197–5201. DOI: <https://doi.org/10.1109/ICASSP.2014.6854594>
- McFee, B., Nieto, O., Farbood, M. M., and Bello, J. P.** (2017). Evaluating hierarchical structure in music annotations. *Frontiers in Psychology*, 8:1337. DOI: <https://doi.org/10.3389/fpsyg.2017.01337>
- Nieto, O. and Bello, J. P.** (2016). Systematic exploration of computational music structure research. In *International*

- Society for Music Information Retrieval Conference (ISMIR)*, pages 547–553.
- Nieto, O., Mysore, G. J., Wang, C.-I., Smith, J. B., Schluter, J., Grill, T., and McFee, B.** (2020). Audiobased music structure analysis: Current trends, open challenges, and applications. *Transactions of the International Society for Music Information Retrieval*, 3(1). DOI: <https://doi.org/10.5334/tismir.78>
- Ong, B. S. and Herrera, P.** (2005). Semantic segmentation of music audio. In *Proceedings of the International Computer Music Conference*, page 61.
- Oyama, T., Ishizuka, R., and Yoshii, K.** (2021). Phaseaware joint beat and downbeat estimation based on periodicity of metrical structure. In *International Society for Music Information Retrieval Conference (ISMIR)*, pages 493–499.
- Paulus, J., Müller, M., and Klapuri, A.** (2010). State of the art report: Audio-based music structure analysis. In *International Society for Music Information Retrieval Conference (ISMIR)*, pages 625–636.
- Raffel, C., McFee, B., Humphrey, E. J., Salamon, J., Nieto, O., Liang, D., and Ellis, D. P. W.** (2014). mir_eval: A transparent implementation of common MIR metrics. In *International Society for Music Information Retrieval Conference (ISMIR)*, pages 367–372.
- Salamon, J., Nieto, O., and Bryan, N. J.** (2021). Deep embeddings and section fusion improve music segmentation. *IEEE Signal Processing Letters*, 24(3):279–283. DOI: <https://doi.org/10.1109/LSP.2017.2657381>
- Sargent, G., Bimbot, F., and Vincent, E.** (2016). Estimating the structural segmentation of popular music pieces under regularity constraints. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 25(2):344–358. DOI: <https://doi.org/10.1109/TASLP.2016.2635031>
- Serrà, J., Müller, M., Grosche, P., and Arcos, J. L.** (2014). Unsupervised music structure annotation by time series structure features and segment similarity. *IEEE Transactions on Multimedia*, 16(5):1229–1240. DOI: <https://doi.org/10.1109/TMM.2014.2310701>
- Shiu, Y., Jeong, H., and Kuo, C.-C. J.** (2006). Similarity matrix processing for music structure analysis. In *Proceedings of the 1st ACM Workshop on Audio and Music Computing Multimedia*, pages 69–76. DOI: <https://doi.org/10.1145/1178723.1178734>
- Smith, J. B., Burgoyne, J. A., Fujinaga, I., De Roure, D., and Downie, J. S.** (2011). Design and creation of a large-scale database of structural annotations. In *International Society for Music Information Retrieval Conference (ISMIR)*, pages 555–560.
- Turnbull, D., Lanckriet, G. R., Pampalk, E., and Goto, M.** (2007). A supervised approach for detecting boundaries in music using difference features and boosting. In *International Conference on Music Information Retrieval (ISMIR)*, pages 51–54.
- Ullrich, K., Schluter, J., and Grill, T.** (2014). Boundary detection in music structure analysis using convolutional neural networks. In *International Society for Music Information Retrieval Conference (ISMIR)*, pages 417–422.
- Wang, J.-C., Smith, J. B., Lu, W.-T., and Song, X.** (2021). Supervised metric learning for music structure features. In *International Society for Music Information Retrieval Conference (ISMIR)*, pages 730–737.

TO CITE THIS ARTICLE:

Marmoret, A., Cohen, J. E., and Bimbot, F. (2023). Barwise Music Structure Analysis with the Correlation Block-Matching Segmentation Algorithm. *Transactions of the International Society for Music Information Retrieval*, 6(1), 167–185. DOI: <https://doi.org/10.5334/tismir.167>

Submitted: 30 March 2023 **Accepted:** 02 November 2023 **Published:** 30 November 2023

COPYRIGHT:

© 2023 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <http://creativecommons.org/licenses/by/4.0/>.

Transactions of the International Society for Music Information Retrieval is a peer-reviewed open access journal published by Ubiquity Press.