

Persistent Homology on Musical Bars

Victoria Callet^[0000–0003–1501–5670]

IRMA, UMR 7501, CNRS, Université de Strasbourg, France
victoria.callet@math.unistra.fr

Abstract. This article presents a new way of building a filtered simplicial complex from a music piece and applying persistent homology in the context of musical analysis. Our approach consists of considering any musical score as the set of its musical bars, which we see as subsets of \mathbb{R}^3 . With this definition, we may consider the Hausdorff distance between two musical bars, which gives us a point cloud from any score, and that allows us to build the associated Vietoris-Rips complex. We will then use barcodes to visualize persistent homology and give an illustration of our construction on a famous movie music piece.

Keywords: Musical bars · Filtered complex · Persistent homology · Barcodes · Musical analysis.

1 Introduction

1.1 Persistent Homology

A filtration of a simplicial complex K is a nested sequence of sub-complexes $\emptyset = K^{-1} \subset K^0 \subset \dots \subset K^N = K$ of K : we call K a **filtered complex**. A simple filtration is presented for instance in figure 1. Starting from a filtered complex, we can compute its simplicial homology (over \mathbb{F}_2) at each **time** of the filtration, and persistent homology gives information about inclusions between the various complexes, as explained in [7]. The associated homology groups $H_*(K^s)$ are characterized by their dimensions, which are called the **Betti numbers**. We can visualize persistent homology on a figure called a **barcode**, where the horizontal axis represents progress in the filtration and a bar that starts at time s and ends at time t corresponds to a generator of $H_*(K^s)$ that is still one for $H_*(K^{t-1})$ but not anymore at time t (see [4]). Barcodes allow us to immediately identify classes that **persist** during the filtration. For instance, barcodes associated to the filtered complex in figure 1 are presented in figure 2, in degrees 0 and 1.




Fig. 1: A filtered complex with 6 times of filtration.



Fig. 2: Barcodes for filtration of figure 1 in degree 0 (left) and degree 1 (right).

Persistent homology is commonly used in a recent field known as **Topological Data Analysis** (TDA), where the main idea is to use topological structures to study and sometimes "recognize" some objects. The general process is summarized in figure 3, where in our context the starting object is a musical score.




Fig. 3: Topological Data Analysis process.

1.2 Context and Problematic

TDA via persistent homology was already used in the context of musical analysis and especially in automatic classification of musical style, as we can see in [5], [2], [3] or [1]. The main and common issue is always the following problematic: **how should we associate a filtered complex with a given musical piece?**

The papers we cited above have different approaches to try to find a consistent answer to this question, using pitch-class sets complexes, time series or the Tonnetz. In our work, we chose to consider that a music piece can be represented by a set of distinct musical bars, which can be thought of as subsets of \mathbb{R}^3 , as defined in section 2.2. Our starting object is a musical score, from which we will extract a point cloud of \mathbb{R}^3 by considering its musical bars and the associated Hausdorff distance, as presented in section 2.3. To build the filtration, we will use the Vietoris-Rips method, which is described in section 2.1. Once we obtain a filtered complex, we can compute persistent homology, i.e. barcodes. In section 3, we will illustrate our approach by giving an analysis of the French movie musical piece: *Comptine d'un autre été: l'Après-midi*, by composer Yann Tiersen. In order to allow musical interpretation, we will be interested in dimension 0 and 1 initially, and we will show that these barcodes, especially in degree 0, can capture structural information about the piece.

2 Persistent Homology on Musical Bars

2.1 Filtration: the Vietoris-Rips Method

The basic object of a Vietoris-Rips complex is a **point cloud**, that is, a set of data points with a metric over it.

Definition 1. Let $X = \{x_1, \dots, x_n\}$ be a point cloud and $\epsilon \geq 0$ be a **parameter**. The **Vietoris-Rips complex** $\mathcal{R}_\epsilon(X)$ is the simplicial complex with X as its set of vertices and $\sigma = \{x_1, \dots, x_k\}$ is a k -simplex if and only if its vertices are pairwise close, that is, when $d(x_i, x_j) \leq \epsilon$ for all pairs x_i, x_j of σ .

Figure 4 shows the classical construction of a Vietoris-Rips complex starting from a given point cloud X and a parameter ϵ . For two given parameters ϵ and ϵ' such as $\epsilon < \epsilon'$, there is the obvious inclusion $\mathcal{R}_\epsilon(X) \hookrightarrow \mathcal{R}_{\epsilon'}(X)$ and so by simply increasing the parameter ϵ , we will get a natural sequence of nested simplicial complexes, that is, a filtered complex. For instance, $\mathcal{R}_0(X)$ is the 0-dimensional simplicial complex with n connected components while $\mathcal{R}_\epsilon(X)$ will have only one connected component for ϵ large enough.




Fig. 4: The Vietoris-Rips method.

Remark 1. For all $\epsilon \geq 0$, we have a simplicial complex, but in practice we will discretize time by choosing a finite number of values for ϵ . In this paper, we choose to work with all values of the form $\epsilon = t\rho$, where $t \in \{0, 1, \dots, 100\}$ and ρ is a fixed constant, so all our filtrations and also barcodes will be at the same **scale**, as we will see in section 2.3.

2.2 Musical Bars of a Score

The main idea of our construction is to consider that a **musical score** \mathcal{S} is simply a set of its musical bars: $\mathcal{S} = \{\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_n\}$, where the indexing corresponds to the musical flow. In this context, our definition of a bar is the following one:

Definition 2. A **musical bar** is a finite subset \mathcal{B} of \mathbb{R}^3 where an element of \mathcal{B} is called a **note** characterized by three coordinates:

- the **position**, which refers to its place in the bar
- the **duration**, expressed in beats
- the **pitch**, which is the value of the note in term of its fundamental frequency

Example 1. Here is an example of a musical bar coding. The first two coordinates are determined by using the meter and the third is coded in *midicent*, which is related to the position of the note in a piano. Notice that if there are some rests in the bar, we can ignore them because the information is already encoded in the position of the next note.



$$\mathcal{B} = \{(0, 1/2, 71), (1, 2, 69), (3, 1, 72)\}$$

2.3 The score as a Point Cloud

If X and Y are two non-empty subsets of a metric space (M, d) , it is possible to define a metric d_H called the **Hausdorff distance** between X and Y by :

$$d_H(X, Y) = \max \left\{ \sup_{x \in X} d(x, Y) ; \sup_{y \in Y} d(X, y) \right\}$$

with $d(x, Y) = \inf_{y \in Y} d(x, y)$. An illustration of this metric is given in figure 5.




Fig. 5: Calculation of the Hausdorff distance between two metric spaces X (the ellipse) and Y (the square).

In our case, each musical bar of a score $\mathcal{S} = \{\mathcal{B}_1, \dots, \mathcal{B}_n\}$ is a subset of \mathbb{R}^3 so we may naturally consider the Hausdorff distance between \mathcal{B}_i and \mathcal{B}_j for any i, j .

Definition 3. Let \mathcal{B}_i and \mathcal{B}_j be two musical bars. The Hausdorff distance d_H between \mathcal{B}_i and \mathcal{B}_j is defined by

$$d_H(\mathcal{B}_i, \mathcal{B}_j) = \max \left\{ \max_{n_i \in \mathcal{B}_i} \min_{n_j \in \mathcal{B}_j} d_1(n_i, n_j) ; \max_{n_j \in \mathcal{B}_j} \min_{n_i \in \mathcal{B}_i} d_1(n_i, n_j) \right\}$$

where $d_1(x, y) = \|x - y\|_1 = \sum_i |x_i - y_i|$.

Following up on remark 1, we will consider the maximal distance d_{\max} of all distances for a given score $\mathcal{S} = \{\mathcal{B}_1, \dots, \mathcal{B}_n\}$ and define $\rho = \frac{d_{\max}}{100}$ as the **precision** we want to work with. We then consider for each $t \in \{0, 1, \dots, 100\}$ the associated Vietoris-Rips complex $\mathcal{R}_{t\rho}(\mathcal{S})$. Furthermore, instead of speaking about "time t of the filtration", we will now say that we look at the filtration with an **error margin** of $t\%$. Indeed, we think of the presence of an edge between two musical bars as an indication that they are "similar", and the parameter t controls the way in which we choose to make this rigorous; for a small value of t , there are few edges which means that bars are finely distinguished, while for t large enough we allow coarser identifications.

We now have a point cloud from any musical score, so we have defined a new way to associate a filtered complex with a score by considering the Vietoris-Rips complex of section 2.1. Thus, we are now able to compute persistent homology (barcodes) and the next section shows an example of a musical analysis using this approach. Recall the filtration and barcodes from figures 1 and 2 respectively; they correspond to the method we just defined applied to the little score with 4 distinct bars of figure 6, which was built arbitrarily to test it.



Fig. 6: The score from which we build filtration and barcodes of figures 1 and 2.

3 Application: Analysis of a Musical Piece

The idea is now to test our approach on a music piece to see in which way it is useful in the context of musical analysis. As we have already mentioned in the introduction, we are interested in the barcodes in degree 0 and 1. For computing homologies, Betti numbers and barcodes, we used our own algorithms and programs written with the help of the *SageMath* system (based on Python).

The musical piece we choose to analyse is taken from the soundtrack of the French movie *Le Fabuleux Destin d'Amélie Poulain*, directed by Jean-Pierre Jeunet (2001). The music is the famous *Comptine d'un autre été: l'Après-midi* for piano, composed and played by the minimalist composer Yann Tiersen (2001). The version we took for our analysis is extracted from [6].

The score has 53 musical bars, some of which are repeated so it only contains 39 distinct bars, and it is split in two parts: actually, the music has 3 different themes in the first part that are played again one octave higher in the second one. These 3 themes and their respective structures are presented in figure 7. Notice that there are some repetitions of these themes in the original score, but here we suppress them in order to work with distinct musical bars only. Moreover, all the melody is constructed over 4 musical bars that are repeated and which constitute the musical accompaniment from figure 8.



Fig. 7: Part of each theme of *Comptine d'un autre été: l'Après-midi*. The first one goes from B_5 to B_8 and is repeated one octave higher from B_{22} to B_{25} . The second goes from B_9 to B_{12} , then is repeated with extra notes from B_{13} to B_{16} and the whole 8 bars are played one octave higher from B_{26} to B_{33} . The third one goes from B_{17} to B_{21} and is repeated one octave higher from B_{34} to B_{38} , with a slight change in B_{38} to bring us to the end of the piece.



Fig. 8: The musical accompaniment of *Comptine d'un autre été: l'Après-midi*. These 4 bars consists of 4 arpeggiated chords $Em - G - B - D$ and are played once at the beginning of the score without any melody, from \mathcal{B}_1 to \mathcal{B}_4 .




Fig. 9: Barcodes for *Comptine d'un autre été : L'Après-midi* in degree 0 (left) and degree 1 (right).

Let us look at the barcode in degree 0 from figure 9: there are several levels of analysis depending of the error margin we choose to take, and the main idea of persistent homology is to focus on the largest bars (those which *persist*), while the smallest ones can be considered as *noise*. In our case, there are 2 bars that stand out when we take an error margin larger than 21%, that means that the corresponding complex has only 2 connected components. One of them corresponds to the last musical bar \mathcal{B}_{39} of the score, which only consists of the final chord Em played with whole notes, and the other is a large dimensional complex where all the musical bars are connected together. This first analysis shows that the barcode in degree 0 separates the end from the rest of the piece, which is a start. For $t\%$ with $t \leq 8$, there are only small bars so we ignore them as noise. Between 8% and 21%, there are 5, 6 or 7 classes that seem to last and more precisely, we found that with an error margin of 14%, the associated complex in the filtration looks like in figure 10, which is really remarkable: actually, there are 6 connected components and if we look at the vertices, we see that each one corresponds to a theme of the song, except for \mathcal{B}_8 and \mathcal{B}_{28} which have a slight different structure than the rest of the first theme.




Fig. 10: The associated complex of *Comptine d'un autre été: l'Après-midi* with an error margin of 14%: each component characterizes a theme of the piece.

We may then conclude that, with a well-chosen error margin, the barcode in degree 0 captures the structure of this piece by separating its different themes in the associated complex.

On the other hand, barcode in degree 1 displays 3 different one-dimensional cycles that are presented in figure 11. Note that some edges of these cycles linked musical bars of one given theme to the same one octave higher, but not systematically and for now we are not able to interpret these cycles musically.




Fig. 11: One-dimensional cycles from *Comptine d'un autre été : L'Après-midi* that appear with an error margin of 8% (left) and 26% (middle and right).

4 Conclusion and prospect

This paper has presented a new method to provide a filtered complex associated with a musical piece. Moreover, our construction reveals interesting results on the piece we chose to study. Actually, it shows that there is a way to capture the global structure of the piece by using barcodes in degree 0. In contrast, barcodes in degree 1 did not display an obvious musical interpretation, and we plan to focus on this dimension in our future work. In fact, one-dimensional cycles could be related to repeating patterns or musical loops in the score, and we are working on highlighting this interpretation from our construction. We also plan to apply our approach to a more general and diverse corpus of music data to see in which way it can capture the global structure of a music piece.

References

1. Bergomi, M.G., Baratè, A.: Homological persistence in time series: an application to music classification. *J. Math. Music* **14**(2), 204–221 (2020)
2. Bergomi, M.G., Baratè, A., Di Fabio, B.: Towards a topological fingerprint of music. In: *Computational topology in image context*, Lecture Notes in Comput. Sci., vol. 9667, pp. 88–100. Springer, [Cham] (2016)
3. Bigo, L., Andreatta, M.: Filtration of pitch-class sets complexes. In: *Mathematics and computation in music*, Lecture Notes in Comput. Sci., vol. 11502, pp. 213–226. Springer, Cham (2019)
4. Ghrist, R.: Barcodes: the persistent topology of data. *Bull. Amer. Math. Soc. (N.S.)* **45**(1), 61–75 (2008)
5. Sethares, W.A., Budney, R.: Topology of musical data. *J. Math. Music* **8**(1), 73–92 (2014)
6. Tiersen, Y.: *Comptine d'un autre été : L'après midi*. In: Yann Tiersen, *Six pièces pour piano, Vol.2*, pp. 5–7. Id Music (2004)
7. Zomorodian, A., Carlsson, G.: Computing persistent homology. *Discrete Comput. Geom.* **33**(2), 249–274 (2005)