



HAL
open science

The Inconstancy of Music

Florence Levé, Gianluca Micchi, Jean-Paul Allouche

► **To cite this version:**

Florence Levé, Gianluca Micchi, Jean-Paul Allouche. The Inconstancy of Music. Journal of Mathematics and Music, 2023, 17 (1), 10.1080/17459737.2022.2068687 . hal-03319124

HAL Id: hal-03319124

<https://hal.science/hal-03319124v1>

Submitted on 6 Feb 2025

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

The inconstancy of music

Florence Levé^{a,b*}, Gianluca Micchi^b, and Jean-Paul Allouche^c

^a*Université de Picardie Jules Verne, MIS, Amiens, France;*

^b*Université de Lille, CNRS, Centrale Lille, UMR 9189 CRISTAL, F-59000 Lille, France;*

^c*CNRS, IMJ-PRG, Sorbonne, Paris, France*

()

Dedicated to the memory of Laurence Maillard-Teyssier (1975–2020)

*Qui de nous
A pu
S'affranchir de l'absence ?
(Eugène Guillevic)*

A melody is often described as a line of music that evolves through time and, therefore, it is possible to draw its 2D pitch-time representation as a series of points implicitly defining a curve. We introduce to computational musicology a descriptor of this music curve: the *inconstancy*, a function that gives information on the curve's smoothness as well as some of its topological properties. A mathematical analysis of the inconstancy of music is provided, followed by a lengthy application of inconstancy to musicological tasks. We compare the inconstancy of melodic lines with that of typical accompaniment patterns such as the Alberti bass; this analysis, together with the case study of W.A. Mozart's Variations on *Ah! vous dirai-je, maman*, suggests a significant difference in the value of the inconstancy of a music line depending on its function. The inconstancy seems to be correlated also with the compositional style: the analysis on almost 10,000 musical themes of the common practice repertoire shows that Baroque music has higher inconstancy. Finally, we also define a windowed version of the inconstancy for studying longer scores and show the insights one can gain into, for example, structural analysis and cadence detection.

1. Introduction

The computational analysis of music is moving steadily towards an ever-increasing use of machine learning techniques, which have been applied to several different tasks including, but not limited to, structure analysis (Ullrich, Schlüter, and Grill 2014; Buccoli et al. 2016), chord analysis (for a survey, see Pauwels et al. 2019), and functional harmonic analysis (Chen and Su 2018, 2019; Micchi, Gotham, and Giraud 2020). These techniques proved themselves very useful whenever the performance of rule-based algorithms was not satisfactory. For example, in the domain of functional harmonic analysis, tools such as fully-connected neural networks, mixed convolutional-recurrent networks, and transformer models have been proven to efficiently deal with non-chord tones (Ju et al. 2017), key detection (Nápoles López et al. 2020) and Roman numeral annotation

*Corresponding author. Email: florence.leve@u-picardie.fr

([Micchi, Gotham, and Giraud 2020](#)).

Differently from the audio or visual domain, however, the domain of symbolic music often needs to rely on a pre-processed representation of the input data, sometimes helped by carefully engineered features. This is because symbolic data is too scarce to train a completely automatic approach. Certainly, one can put special attention when designing the model: compact, highly-optimised architectures can reduce the total number of weights in the network, effectively reducing the need for data. Also, careful data sampling for the training, validation, and test sets can reduce the bias inherently present in all datasets and contribute to increasing the performance of models without increasing the quantity of collected data ([Huron 2013](#)). Careful feature engineering, however, reducing the complexity of the network in favour of an initial, controllable and deterministic analysis of the data inspired by expert knowledge, is one of the most common tools used to reduce the difficulty of the task for the machine by limiting the scope of its investigation. It also helps with the interpretation of the results: despite recent efforts on this subject, machine learning models trained on raw data are notoriously hard to interpret (see for example [Mishra, Sturm, and Dixon 2017](#)). On the contrary, the integration of musical knowledge can help to evaluate the results in a more qualitative way ([Carsault, Nika, and Esling 2018](#)), thus the use of features derived from music theory can bridge the gap between mathematical or computational analysis and musical interpretation. For all these reasons, there is still a lot of space for sensibly chosen musicological features as the first step of the analysis algorithm, be they rule-based or statistical in nature ([McKay, Cumming, and Fujinaga 2018](#)).

The features thus engineered need not be particularly complicated, as testified by [Rydén \(2020\)](#). Here, the author takes a set of fugue subjects coming from different composers and styles and analyses them according to only six global features: number of notes, pitch range, unique pitch classes, initial interval, unique intervals, and largest interval.

[Bigo et al. \(2018\)](#) instead propose a broad set of 44 binary features that can be derived from symbolic scores and span primarily rhythm and voice leading or harmonic characteristics: for example, rhythm-related features check whether the cadence happens on a strong beat, or whether it is followed by rests, and voice-related or harmonic features include whether the bass has a V-I leap or whether the arrival chord is a perfect triad. All these features can be described as low-level because they are local (they take into account a maximum of 3 nearby notes at a time) and they require simple arithmetic to be derived. These features are chosen due to their expected importance for the task of cadence detection; a different task, therefore, requires a different set of features. This is apparent when considering a subsequent paper published by the same team in which the authors analyse the sonata structure ([Allegraud et al. 2019](#)). In that paper, they use a different set of 26 binary features: tonality, presence of a repeated pattern, presence of a cadence, etc. These features are of a much higher level compared to the ones in the previous work, that is, they require a longer procedure to be derived starting from the score. This is a standard pattern: a complex musicological result is more strongly connected to higher-level features than to the basic ones, and the long-range global structure of the composition is no exception to that.

Another example of how a complex task requires a set of higher-order features is given by [Kirlin and Yust \(2016\)](#) for the task of Schenkerian analysis. In their paper, Kirlin and Yust use 18 features derived from melodic, harmonic, metrical, and temporal analysis of the music as input to a random forest tasked with making a Schenkerian reduction of any given score.

In his book *A Geometry of Music*, [Tymoczko \(2011\)](#) uses a series of symmetry invari-

ances to derive a representation of music that makes some of the important features more apparent.

In this paper, we propose to use a quantity called inconstancy (see e.g. [Allouche and Maillard-Teyssier 2011](#)). The inconstancy of a curve can be defined as the average number of intersection points of a curve with random straight lines. Essentially, if the curve is very similar to a straight line, its inconstancy is low, while the inconstancy of a complex “disheveled” curve is large. It turns out that the inconstancy of a curve is easy to compute thanks to a theorem of Cauchy-Crofton, which gives its value in terms of the length of the curve and the perimeter of its convex hull. [Allouche and Maillard-Teyssier \(2011\)](#) indicated that this “measure of disorder” could possibly be used in biology or in other fields. They ask, “are ‘large fluctuations’ of the Body Mass Index risk factors for cardiovascular diseases in relation with the metabolic syndrome?”, or propose the use of this measure for “analyzing fluctuations of the stock market”, or even for “quantifying the ‘smoothness’ of musical themes” (p. 2269). To explore this last question, our paper proposes, in a vein that might be reminiscent of Tymoczko’s book (see e.g. his Chapter 3, Geometry of chords), to introduce the inconstancy function ([Allouche and Maillard-Teyssier 2011](#)) as a feature extractor for monophonic music. Our goal is to provide a measure that is easy to compute and that can provide some relevant musical information.

We believe that such a quantitative approach can be applied to a variety of situations, including but not limited to: detection of musical structure, for example for the identification of patterns; characterization of melodic lines, by means of contour analysis (i.e. the sequence of pitch intervals shaping the melody) or simply in opposition to typical accompaniment patterns; and generation of music, by use of an appropriate loss function, to produce more stylistically accurate melodies.

The article is structured as follows. In Section 2 we introduce the concept of inconstancy in music and discuss its mathematical properties, with particular attention to invariant quantities. Section 3 shows the typical values of inconstancy on a series of small examples, ranging from typical patterns such as scales and Alberti bass lines to simple melodies taken from folk tunes. Section 4, instead, is a quick overview of possible usages of the inconstancy function in a larger musicological scope: we define the inconstancy on sliding windows of music and show its result on entire pieces from the Western common practice repertoire.

2. The inconstancy of musical lines

2.1. *The inconstancy*

The inconstancy of a curve γ in a two-dimensional space is defined in two different and equivalent ways ([Allouche and Maillard-Teyssier 2011](#)): (1) The average number of points at which a random intersecting line touches or crosses the curve. (2) Twice the ratio between the length of the curve $l(\gamma)$ and the perimeter of its convex hull $p(\text{Hull}(\gamma))$,

$$\mathcal{I}(\gamma) = \frac{2l(\gamma)}{p(\text{Hull}(\gamma))}. \quad (1)$$

The inconstancy of a segment joining two points always equals 1. This can be shown independently using both definitions. Considering definition (1), this result follows from

the fact that two straight lines can only intersect once.¹ Using definition (2), the convex hull is simply the “closed” polygon that goes from one point to the other and back, so its perimeter is exactly twice the length of the curve. We use this result to extend by continuity the definition of the inconstancy on ever-shrinking segments and say that the inconstancy of a curve γ made of a single point is equal to 1 (which is coherent with the first definition).

The inconstancy of a curve γ joining N points always satisfies $1 \leq \mathcal{I}(\gamma) \leq N - 1$: since the curve is made of $N - 1$ segments no straight line can intersect the curve in more than $N - 1$ points. As a consequence, it is easy to show that the inconstancy is a “non-additive” function in the sense that the inconstancy of the concatenation of two sequences is not equal to the sum of the individual inconstancies. Indeed, let us consider a curve made of two broken segments: while each segment has separately an inconstancy of 1, the inconstancy of the total curve can take every value between 1 and 2, depending on the angle at which the two segments join.

A typical use of the inconstancy is to study the property of a sequence $(a_n)_{\mathbb{N}}$. By defining an ordered set of points $S = \{(n; a_n) \mid n \in \mathbb{N}\}$, one can define the total inconstancy of the sequence $\mathcal{I}(S)$ to be the inconstancy of the curve γ joining each point of S with its successor. It is also possible to study the evolution of the inconstancy along the sequence by taking the inconstancy of all subsets containing only the first i points of S , with i going from 1 to the total length of the sequence.

2.2. Inconstancy applied to music

Music is a very high-dimensional object. Even when analysing symbolic music — therefore discarding all information related to performance, instruments, and recording environment — one needs lots of independent dimensions to fully describe a score: pitches, time, dynamics, articulation, tempo, instruments, etc.

From an artistic point of view, each of these dimensions is necessary. However, a lot of musicologically relevant work has been done by representing symbolic music as a set of points² in the standard piano roll notation, which stores only two dimensions: pitch and time. To make a stronger analogy with number sequences, we use an adapted version of the piano roll notation in which note durations are discarded. Therefore, each note is represented as a point-like event at a given offset from the beginning of the score. Also, *we will only focus on monophonic data* and discard occasional polyphony when analysing musical scores (see Section 3.4.1 for a discussion on how we extract monophonic music from polyphonic scores).

Given a musical score S made of a succession of N notes $n_i = (t_i; p_i)$, each described by a time offset t_i and a pitch p_i , we can define the line γ_S that follows the score as the concatenation of all segments joining pairs of successive points. The inconstancy of the musical score S is then simply the inconstancy of γ_S . In contrast to [Allouche and Maillard-Teyssier \(2011\)](#), we do not consider the origin of the axes as the first point of the line.

As an example, we plot in [Figure 1](#) the evolution of the inconstancy along the musical line for the beginning of the French folk song *Ah! vous dirai-je, maman*.³

On the horizontal axis, we have plotted the musical time t (in measures) elapsed since

¹Unless they are the same line in which there is a problem of defining intersections properly, but this case has infinitesimal weight and can be discarded.

²Or better, segments, but in the following we will ignore note durations and focus only on their onset.

³We will go back to this melody in [Section 3.4](#), where we analyse the inconstancy of the 12 variations that Mozart wrote on this theme.

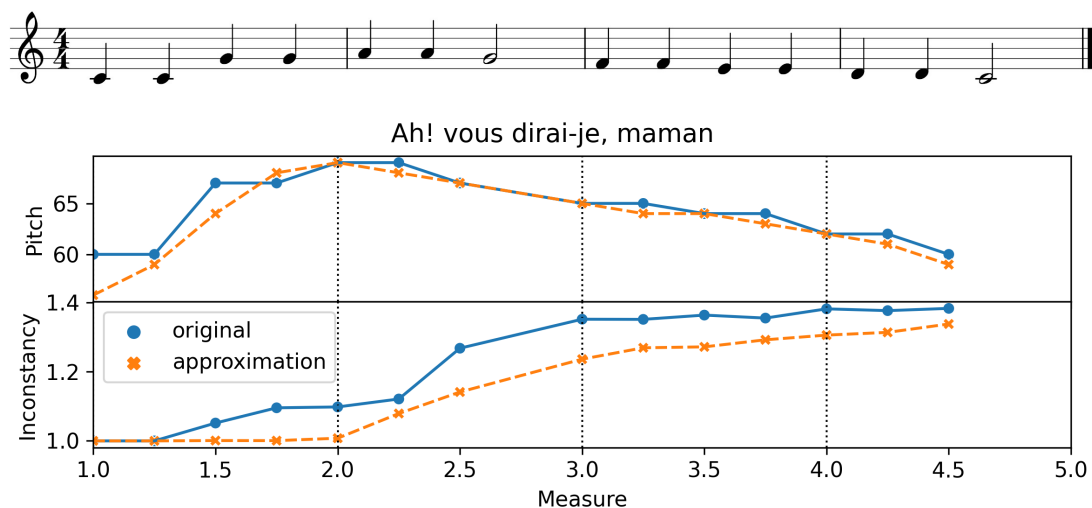


Figure 1. Top: the first four measures from the French folk song *Ah! vous dirai-je, maman*. Middle: the representation of the score as a sequence. Bottom: inconstancy $\mathcal{I}(\gamma_S(t))$ for the portion of the melody comprised between time 0 and t (in measures). Dashed lines correspond to the polynomial approximation of the melody (see details in the text); we sample the value of the approximation at the same times as the points in the original representation and quantize the MIDI pitch number to the closest integer.

the beginning of the melody, while the vertical axis represents respectively the MIDI pitch of the notes (top) and the value of the inconstancy of the melodic line up to offset time t (bottom). Broken lines have been drawn to join the discrete values and help reading.

We can see two main features: first, the value of the inconstancy has a big jump at measure $t = 2.5$, when the direction of the music is reversed for the first time (the G descends from the A, instead of keeping the ascending movement); second, the inconstancy in the second half of the musical phrase does not change much in value. This is a combined effect of the increased length of the score, as we will see in Section 2.6, and of the general structure of the music, which resembles a straight segment (cf. top panel).

The inconstancy, therefore, is associated with the geometrical shape of the contour of the melody. Melodic contour has been the subject of lots of research over time, for example, the seminal work of [Huron \(1996\)](#), where melodies are classified according to the general shape of the contour. We believe that inconstancy could help the classification algorithm by providing an additional quantity to study since melodies with similar global shapes share similar values of inconstancy. As an example of that, we calculated the inconstancy of the polynomial approximation of *Ah! vous dirai-je, maman* (see Figure 1). This approximation was introduced in [Müllensiefen and Wiggins \(2011\)](#) as the low-order polynomial that minimizes the residues with the point-like representation of the melody. As expected, the inconstancy of the original and approximated melody have quite similar values, the biggest deviation happening at the beginning where the approximated melody is much smoother.

2.3. Mathematical properties of the inconstancy of music

There is at least one fundamental difference between the inconstancy of a sequence of numbers and the inconstancy of music: music has physical dimensions, therefore units of measure, and the value of the inconstancy depends on the unit of measure chosen. For

example, the time axis can be expressed in quarter lengths, but also in eighth lengths, measures, beats, seconds; while the pitch axis is typically expressed with MIDI numbers, but other admissible choices are, for example, the interval with respect to the first note of the score and the fundamental frequency of each pitch.

It is therefore of the utmost importance to study how the inconstancy behaves under the manipulation of the axes. Let us explore some of these properties.

2.3.1. Behaviour under translation

Since the result of Equation (1) is uniquely determined by a ratio of the sum of segment lengths, it is preserved under transformations that preserve segment lengths such as translations. This means, for example, that transposing the music to a different key does not change its inconstancy.

2.3.2. Behaviour under orthogonal transformations

The inconstancy of a curve is invariant under orthogonal transformations such as rotations or reflections, too. This follows from the fact that orthogonal transformations preserve the inner product between two vectors and, therefore, also segment lengths. Musically speaking, this means that contrapuntal techniques such as the vertical inversion (all ascendant intervals become descendant) or the reversal (reading the music from the right to the left) do not affect the inconstancy value. It is maybe useful to point out that the intermediate values of the inconstancy along a musical line change when considering the original melody and its reversed version, but only because the notes kept are different.

2.3.3. Behaviour under isotropic scaling

Let us consider a transformation $T_\lambda : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ such that $T_\lambda(x, y) = (\lambda x, \lambda y)$ with $\lambda \in \mathbb{R}$, $\lambda \neq 0$. Let us define by $T_\lambda \gamma$ the application of the transformation T_λ to the space where γ is defined. Then $\mathcal{I}(\gamma) = \mathcal{I}(T_\lambda \gamma)$, since both the numerator and the denominator of Equation (1) would be multiplied by a factor $|\lambda|$. This transformation has no direct musical interpretation, but we will use it to determine what happens changing unit of measure.

These three operations (translation, orthogonal transformation, and isotropic scaling) are not sufficient to completely determine the behaviour of the inconstancy under an arbitrary affine transformation in the 2D plane⁴. We still need a fourth, independent operation.

2.4. Behaviour under stretching of one axis

This is probably the most important property of them all. It determines how the inconstancy varies under a linear change in the unit of measure. For example, changing a quarter note to a half note is equivalent to stretching the time axis by a factor of 2, changing it to an eighth note is equivalent to stretching the time axis by a factor 0.5.

Throughout this section we will refer to the application of the inconstancy to the theme of Chopin's Nocturne Op. 9 No. 1 in Figure 2.

⁴We leave the study of non-linear transformation such as a conversion on the pitch axis from musical pitches to fundamental frequency (exponential relation) to a future work.

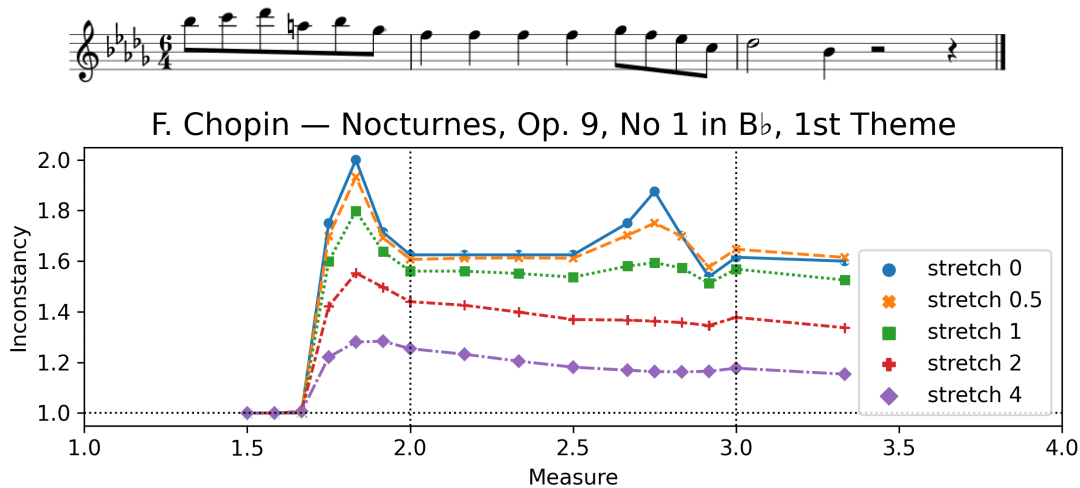


Figure 2. The first theme of Chopin’s Nocturne Op. 9 No. 1 in Bb minor. Each of the five lines corresponds to a different value of the stretch coefficient, where stretch 1 means no stretch applied i.e. standard inconstancy, and stretch 0 corresponds to projected inconstancy.

We define an operation S_λ so that $S_\lambda(x, y) = (\lambda x, y)$, with the constraint $\lambda > 0$. Let γ be, as usual, the curve joining a set of points (x_i, y_i) . We also define $S_\lambda \gamma$ as the curve obtained by joining all the points $S_\lambda(x_i, y_i)$. We call S_λ a *stretch*. It is sufficient to study the stretch along the time axis since a stretch S'_λ along the pitch axis can be obtained, for example, by means of stretch $S_{1/\lambda}$ on the time axis followed by a global scaling T_λ .

Let us start by finding what happens to the inconstancy when the parameter λ tends to 0 or to $+\infty$. We recall Equation (1), which says that the inconstancy of a curve γ is the ratio between twice its length and the perimeter of its convex hull.

2.4.1. Limit when $\lambda \rightarrow 0$

In the case $\lambda \rightarrow 0$, the curve $S_\lambda \gamma$ tends to a succession of a finite number of (possibly overlapping) vertical segments. Therefore, the convex hull surrounding $S_\lambda \gamma$ is the degenerate rectangle with zero-length horizontal edges and whose vertical edges both coincide with the segment between y_{\min} and y_{\max} . The perimeter of the hull is, therefore, $p = 2(y_{\max} - y_{\min}) \equiv 2\Delta y$, also known as the *ambitus* of music. The length of the curve, instead, is given by the sum of the length of all the (possibly overlapping) vertical segments, $l = \sum_{i=1}^N |y_{n+1} - y_n| \equiv N\bar{y}$, where \bar{y} is the average vertical distance of a segment (or *melodic interval*) in the music. Therefore:

$$\lim_{\lambda \rightarrow 0} \mathcal{I}(S_\lambda \gamma) = \frac{N\bar{y}}{\Delta y}. \tag{2}$$

This limit is used in the calculation of the projected inconstancy (see Section 2.5).

2.4.2. Limit when $\lambda \rightarrow +\infty$.

As a first step, let us apply a global scaling of a factor $1/\lambda$, under which the inconstancy is invariant. After the stretch along the time axis, each point in the line γ , therefore, is subject to the final transformation $T_{\frac{1}{\lambda}} S_\lambda(x, y) = (x, \frac{1}{\lambda} y)$. A discussion similar to the one for the case $\lambda \rightarrow 0$ follows. However, in the case of music, the x -axis encodes the time,

and music is written in a time-ordered manner. This entails that γ is monotonic in x , a condition that implies that no segments making up the compressed line $\lim_{\lambda \rightarrow +\infty} T_{\frac{1}{\lambda}} S_{\lambda} \gamma$ overlap.⁵ Therefore the total length of the curve is exactly half the perimeter of its hull and the result for the inconstancy is

$$\lim_{\lambda \rightarrow +\infty} \mathcal{I}(S_{\lambda} \gamma) = \lim_{\lambda \rightarrow +\infty} \mathcal{I}(T_{\frac{1}{\lambda}} S_{\lambda} \gamma) = 1. \quad (3)$$

Figure 2 shows values for different stretch values. A stretch coefficient of 100 (not included in the figure) gives an inconstancy of 1.0004 when considering the entire Chopin theme.

2.4.3. *What happens in between.*

The inconstancy is a smooth function of the stretch parameter λ . This is due to two considerations:

- (1) the convex hull is defined in such a way that all points of the curve are either on the hull or inside it: a stretch will not change this topological property and each point will keep its relative position with respect to the hull;
- (2) the length of a single segment is a smooth function of λ .

It follows that the inconstancy is the ratio of two smooth functions which are always positive. As a recap, Figure 3 shows the inconstancy as a function of λ (solid lines) and its derivative for the musical example of Figure 2. Besides what was already discussed, we can see that the inconstancy has a small peak around $\lambda \approx 0.3$) and its derivative shows a clear minimum just above the value $\lambda = 1$.

Remark: despite it being continuous in the stretch coefficient λ , the inconstancy is not a continuous function in general. Consider e.g. the following sequence of curves: C_0 is the upper half-circle of radius 1 and center 1; C_1 consists of the two upper half-circles of radii $1/2$ and centers $1/2$ and $3/2$; for C_n we have 2^n upper half-circles of radii $1/2^n$ with centers at $1/2^n, 3/2^n, \dots$. The length of each curve C_n is π , the inconstancy of C_n tends to $\pi/2$, while C_n tends to the segment $[0, 2]$ whose inconstancy is 1.

2.4.4. *Stretched inconstancy.*

In our experiment with the corpus of musical themes by [Barlow and Morgenstern \(1948\)](#), we have calculated the inconstancy on all the themes after applying a stretching along the time axis. We always find the derivative of the inconstancy as a function of the stretch coefficient to show exactly one (very clear) global minimum. The existence of this minimum is almost always guaranteed. This can be shown in three steps: (1) the inconstancy is a continuously differentiable function of the stretch coefficient λ defined for values of λ between zero and positive infinity; (2) on one side, the value of the inconstancy at infinite stretch tends to one, therefore its derivative tends to zero; and (3) on the other side, the value of the inconstancy at zero stretch is finite and greater than or equal to one.

Let us take the case in which the limit when $\lambda \rightarrow 0$ is strictly greater than one. Since the inconstancy is a continuous function of the stretch coefficient, the derivative must take some negative value at some point. Therefore, it follows that it must have at least one minimum, possibly at $\lambda = 0$.

⁵The same is not valid in general for the y -axis: all but the simplest pieces of music have melodic intervals in both directions, and therefore the segments joining them overlap when projected on the y -axis.

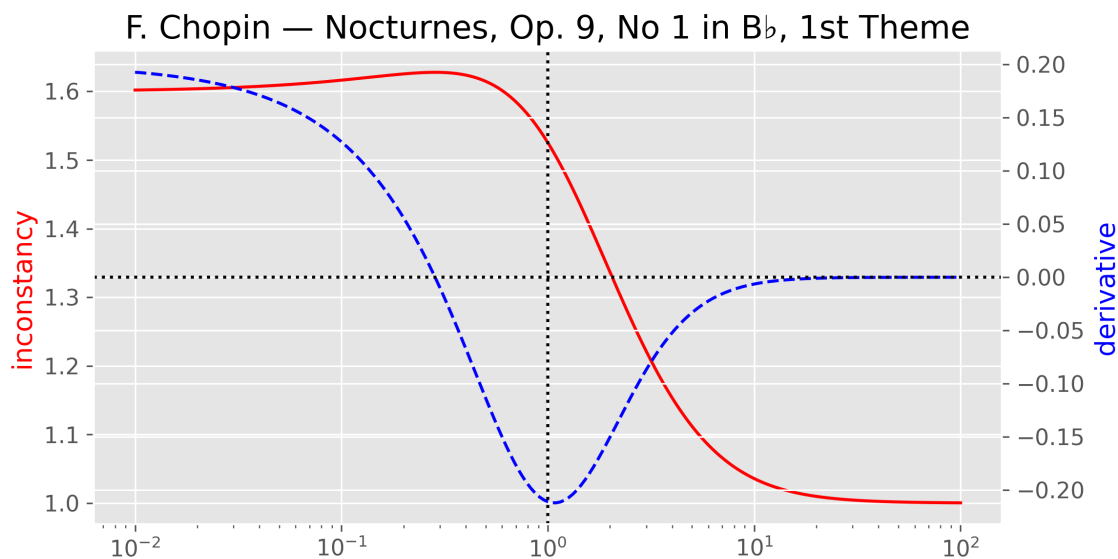


Figure 3. The inconstancy as a function of the stretch coefficient λ (solid line) and its derivative. Notice that the x -axis representing the stretch coefficients is in a logarithmic scale, so the correspondence between the function and its derivative is not obvious at first glance. As a guide, notice that the red line is initially increasing, which is compatible with the blue line being above 0. It is also easy to directly map the position where the sign of the derivative changes with the location of the maximum in the inconstancy value. Due to the logarithmic x axis, it is more difficult to visually identify the position of maximum slope in the red line. The horizontal dotted line corresponds to derivative zero, and the vertical dotted line to $\lambda = 1$, that is, no stretch applied.

In the other case, of inconstancy going to one when $\lambda \rightarrow 0$, the stretched inconstancy can either be a constant as a function of λ , in which case we do not have to worry about the choice of the unit of measure⁶; or it can vary. In this latter case, since it cannot get to a value smaller than one, it must have a maximum at some finite value λ_m ; hence, we can apply the same reasoning as before to show that there must be a minimum in the derivative to the right of λ_m .

For reason of mathematical parsimony, we define the time-independent stretch λ_0 as the smallest stretch at which the derivative of the inconstancy with respect to the stretch coefficient λ has a minimum and the *stretched inconstancy* as the inconstancy of the curve $S_{\lambda_0}\gamma$. By construction, the stretched inconstancy does not depend on the choice of the unit of measure and therefore removes one degree of freedom from the problem. As a consequence, it also is invariant with respect to compositional techniques such as repeating a theme in rhythmic diminution or augmentation.

2.5. Projected inconstancy

Let \vec{a} be the unitary vector defining the direction of an axis A . The projection of a point P on axis A is defined in the standard way as

$$P_A = \frac{\vec{p} \cdot \vec{a}}{\|\vec{a}\|} \frac{\vec{a}}{\|\vec{a}\|},$$

⁶As an example of a melody whose inconstancy is constantly equal to 1 for all values of the stretch coefficient, we can take a melody that comprises a single pitch or, more generally, a melody in which the ratio between the melodic and inter-offset intervals is constant.

where \vec{p} is the vector defined by the endpoint P .

Let us define the projection of the curve γ on the axis A as the curve γ_A that joins all the projected points, and the projected inconstancy \mathcal{I}_A as the inconstancy of γ_A . In the special case for which the projection axis is the y -axis, hence $\vec{a} = (0, 1)$, the projected inconstancy \mathcal{I}_y takes the value defined in Eq. 2. This quantity has the desired property of being independent of the chosen units of measure. For the rest of the paper, when we refer to the projected inconstancy without specifying which axis, we always understand the y -axis.

The projected inconstancy measures how much the music folds onto itself, regardless of the time dimension. It is a topological measure of the contour of the melody: a purely ascending (descending) melody has a projected inconstancy of exactly 1, while a purely convex (concave) melody — that is, a melody that ascends and descends (descends and ascends) exactly once, going back to the initial pitch — has a projected inconstancy of exactly 2.

2.6. Inconstancy as a function of the score length

Another fundamental characteristic of music is that, while it can be extended indefinitely on the time axis by simply adding more notes to it, it is always limited in the pitch axis; this limit is determined by the instruments that play that music. For example, in the case of piano music, the maximum ambitus for pitches extends from 21 to 108.

This consideration suggests the study of the behaviour of the inconstancy function when the ratio between the total duration of the music and its total pitch interval goes to infinity in the absence of any stretch. Another way to express this constraint is

$$\lim_{\epsilon \rightarrow 0} \mathcal{I}(\gamma) = ?, \quad \text{with } \epsilon = \frac{\Delta y}{\Delta x}, \quad (4)$$

where Δx and Δy are the support of γ on x and y , respectively; that is, the total duration and the total ambitus of pitches used.

Let us assume that the typical segment joining two successive notes in the curve has components (\bar{x}, \bar{y}) . Let us also define $\eta = \bar{y}/\bar{x}$ the ratio between these typical components. The parameter η is a characteristic of the musical piece and the units of measure chosen.⁷ Since we are working with monophonic music, $\bar{x} = \Delta x/N$ and the typical length of a segment can be written as $\bar{l} = \Delta x \sqrt{1 + \eta^2} / N$. The total length of the curve is therefore $l = N\bar{l} = \Delta x \sqrt{1 + \eta^2}$.

On the other hand, instead, the convex hull of γ is reduced, at first order of approximation in $\epsilon \ll 1$, to a rectangle of sides Δx and Δy . Hence, its perimeter is $s = 2\Delta x(1 + \epsilon)$.

Putting both together, Equation (4) becomes

$$\lim_{\epsilon \rightarrow 0} \mathcal{I}(\gamma) = \lim_{\epsilon \rightarrow 0} \frac{2\Delta x \sqrt{1 + \eta^2}}{2\Delta x(1 + \epsilon)} = \sqrt{1 + \eta^2}. \quad (5)$$

Therefore, the inconstancy tends to a finite value that does not depend on the number of notes N when this goes to infinity. Let us point out how this result differs from the one in Equation (3): a short piece of music that has been infinitely stretched along the

⁷To give an idea, on the Barlow dataset, comprising almost 10,000 musical themes, the average values of \bar{x} and \bar{y} are respectively 0.73 and 2.4, meaning that the average note length is a dotted eighth note and the average distance between two consecutive notes is a little more than a whole tone.

x -axis to enforce the constraint $\epsilon \rightarrow 0$ while keeping a finite number of notes N will also have $\eta \rightarrow 0$, therefore inconstancy 1.

We conclude this section by briefly commenting on what happens to projected and stretched inconstancy in the situation of long themes. For the former, it completely disregards any temporal information and just adds up contributions from all notes. As the music grows longer, its ambitus and average melodic interval tend to stabilize, therefore their ratio becomes constant. Equation (2) then tells us that the projected inconstancy grows linearly with N . The latter also seems to be almost linear in N . What happens is that the time-independent stretch becomes extremely small for long themes so that the stretched inconstancy approaches the projected inconstancy.

3. The inconstancy of musical patterns

This section is devoted to getting an intuitive grasp of inconstancy and its evolution along a musical line. We will see the typical inconstancy of common musical patterns such as scales, arpeggios, and Alberti bass before analysing some simple melodies. The section ends with an analysis of the 12 variations written by Mozart on the French folk-song *Ah! vous dirai-je, maman* (see theme in Figure 1).

3.1. Major scales.

We will start our examination with one basic musical sequence: a major scale, represented in Figure 4.

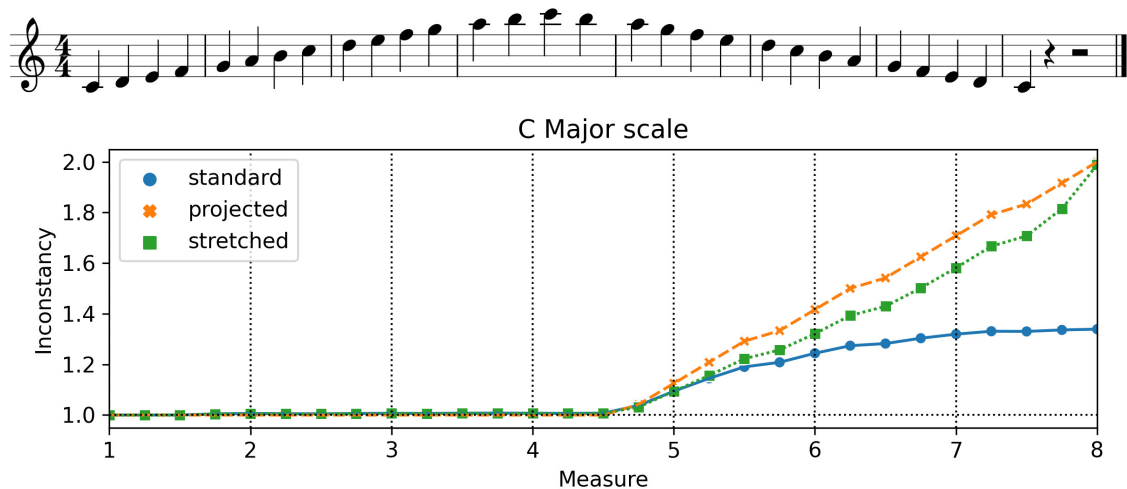


Figure 4. Top: scale of C major. Bottom: evolution of its standard, projected, and stretched inconstancy.

The inconstancy is almost equal to 1 all the way to the highest pitch. This is because the ascending movement is monotonic and can be represented with an almost straight line: the rhythm is steady and the distance between successive pitches is almost constant (often +2 semitones, sometimes +1).

When, starting from the C6 in measure 4 (offset 15), the direction is reversed, the inconstancy begins to steadily increase. But while the standard inconstancy gets to a plateau quite fast, the projected inconstancy increases until reaching a value of 2, as was expected (see Sec. 2.5). This is a confirmation that the inconstancy can be used

for a simple topological investigation of the properties of the sequence. The stretched inconstancy follows quite closely the projected inconstancy. This is also quite intuitive: the representation of the score is very similar to the two oblique edges of an isosceles triangle; and this incomplete triangle, whose maximal inconstancy is 2 because it only has 2 edges, can be stretched to two identical vertical segments, one “going up” and the other “coming down”, for which the inconstancy is exactly 2.

3.2. Accompaniment patterns.

Other common musical sequences are the accompaniment patterns that are commonly used to support a melody. In order to analyse their inconstancy, we have written four different simple patterns based on the bassline of Pachelbel’s *Canon in D*, shown in Figure 5 from top to bottom: one with repeated notes, two with different arrangements of the Alberti bass, and one with arpeggios.



J. Pachelbel - Canon in D (accompaniment patterns)

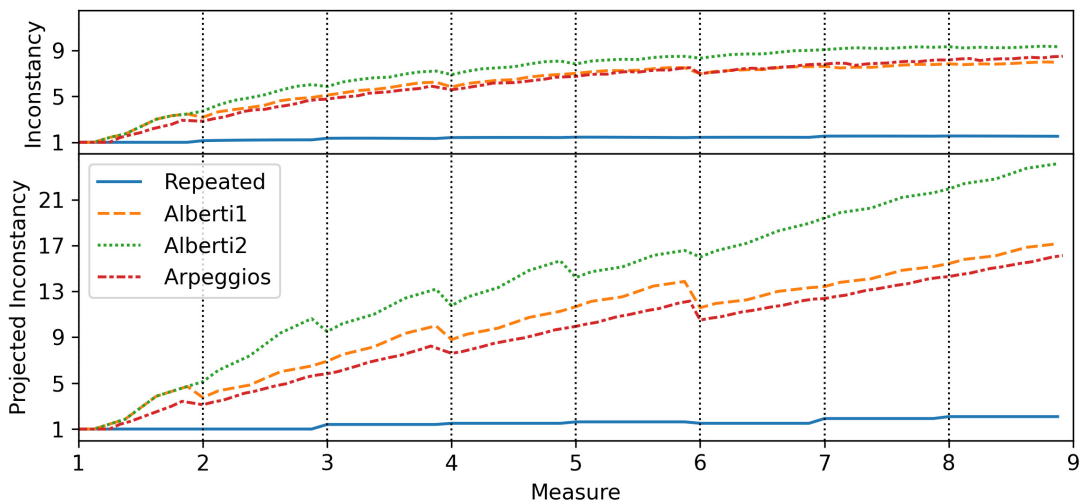


Figure 5. Inconstancy of four different realizations of the accompaniment for Pachelbel’s Canon in D. The vertical lines separate different measures.

The plot shows once again that the inconstancy of long themes tends to saturate, but not the projected inconstancy.

The first feature we focus on is the sudden decrease in value that the inconstancy has at the beginning of each measure in the first half of the score. This decrease is especially apparent in the projected inconstancy. The explanation is quite simple, once one considers Equation (2): at the beginning of each measure, at least for the first six measures, the ambitus increases, therefore making the denominator a larger quantity. Notice, for example, that at the beginning of measure five the dashed orange line corresponding to the first arrangement of the Alberti bass, named Alberti1, does not take a hit: this is because the ambitus of the music has not changed.

Let us now focus on the classification of the patterns. It is very easy to recognize the bass line made only of repeated notes because its inconstancy is very small all through the piece. A very similar inconstancy would also be obtained for an even simpler bass line made only of whole notes.

However, the inconstancy is not enough to tell apart arpeggios in triplets from Alberti bass. The proof of that is in the fact that there is more difference between two different realizations of Alberti bass than between one of those and an arpeggio. This is due to the different use of inversions in the two instances of the Alberti bass. Geometrically, it can be explained in an intuitive manner: if one were to draw a straight line through the entire pattern, one would find that the notes in the second version of the Alberti bass cross the line much more often than in the first line, where they tend to have sensibly lower average value in the middle rather than on the sides (showing, therefore, a global concave pattern).

It is instructive, however, to analyze the difference also from the segment length perspective. Let us do that for the projected inconstancy, which is the simplest of the lot.

The first Alberti bass only uses triads in their root position, which can be described in pitch class notation as $(0, 7, 4, 7)$ or $(0, 7, 3, 7)$ for, respectively, major and minor triads. The total length of the curve inside a single measure, therefore, is always equal to $2 \times (7 + 3 + 3) + 7 = 33$ for major triads and $2 \times (7 + 4 + 4) + 7 = 37$ for minor.

The second Alberti bass, instead, also has chords in their first inversion (at measures 2, 4, and 6). The typical curve length for measures using first inversions is 44 for major triads and 47 for minor.

The average intermeasure segment length, instead, is roughly the same, meaning that the total curve is longer in the second case. Since the total length of the contour is also almost invariant, we conclude that the second Alberti bass has a higher inconstancy.

Finally, let us take a look at the arpeggios. An upper bound for their projected inconstancy value is 32 because there are a total of 31 changes of direction in the melody. Its actual value is 16.125 meaning that the overlap between different measures is just above 50%.

3.3. *A simple melody.*

We pass now to a less academic, albeit still very simple, example: the French folk song *Au clair de la lune*, represented in Figure 6.

Let us focus on the standard inconstancy first (solid lines). The first note for which the inconstancy is different from one is at offset $t = 1.5$, which corresponds to the fourth note in the melody, i.e. the first D. This is understandable because the first three notes are the same, so they generate two perfectly aligned segments which, combined, give inconstancy 1. One can follow the inconstancy on the rest of the melody by comparing it to the score. It increases sharply especially in two sections: between offsets 2 and 3 first, and then again between 4 and 5; these correspond to reversals in the direction of the melody. On the other hand, when repeating notes or keeping the same direction, the inconstancy tends to stay flat or even decrease.

The second section of the music, despite being identical to the first one, shows an evolution of the inconstancy that is much flatter. This is a clear example of the effect that the increasing length of the music has on the value of the inconstancy, as discussed in Section 2.6.

This is no longer true when studying the projected and stretched inconstancy, for which the evolution in the second half is very similar to the evolution in the first half.

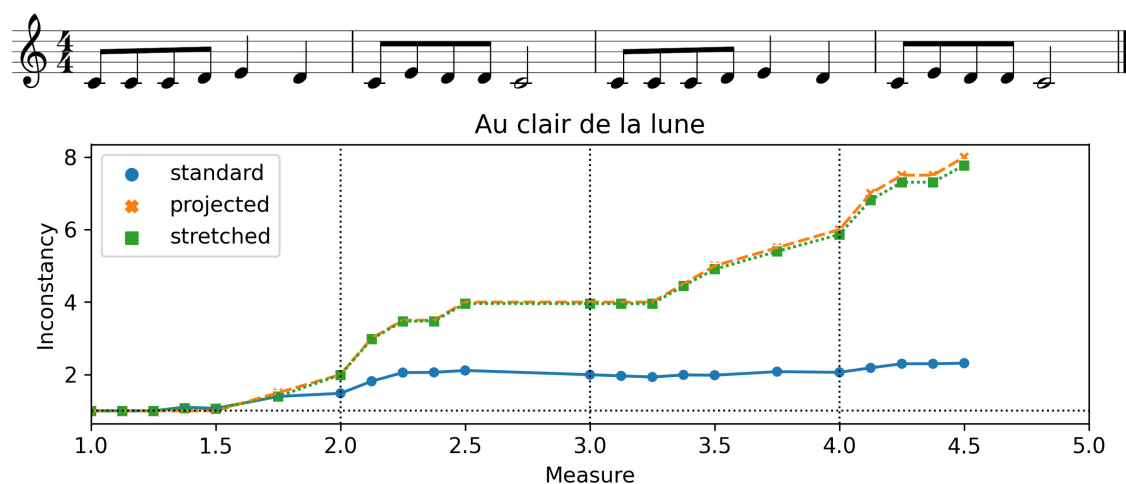


Figure 6. The first four measures of the French folk song *Au clair de la lune*.

3.4. A study on the variations *Ah! vous dirai-je, maman*

We conclude this introductory exploration by analysing the twelve variations that Mozart wrote on the French folk song *Ah! vous dirai-je, maman* (already introduced in Fig. 1).

3.4.1. Monophonic melody extraction

Since these variations are polyphonic, we need to extract monophonic melodies before calculating the inconstancy. We consider left and right hands separately and we apply to each of them the all-mono algorithm presented by [Uitdenbogerd and Zobel \(1999\)](#). For any offset, this algorithm takes the note with the highest pitch. It is therefore also known as the *skyline* algorithm. A known issue with this algorithm is that it may still lead to overlapping tones if a high note is kept for a long duration, while other lower notes are played. To avoid this issue, we propose two versions of this algorithm. The first one, the original one, stops the current note when a new note enters; the second one keeps the current note for its entire duration, preventing other notes from being written before the current one ends. Those two versions are exemplified in Fig. 7.

We can see that the choice of the algorithm makes a big difference, both in terms of the music kept and in terms of the inconstancy. It is outside the scope of this paper to discuss the best algorithm for monophonic extraction of melodies, but we want to show that the inconstancy is definitely influenced by the chosen algorithm. With the opposite perspective, the inconstancy could be used to direct the best choice for monophonic melody extraction.

3.4.2. Their inconstancy

In Figure 8 we study the inconstancy of the first phrase (eight measures) from the theme and each of the twelve variations that Mozart wrote for this folk melody. We apply the modified version of the skyline algorithm to retrieve monophonic music.

All variations are written in 2/4 except for variation 12, which is written in 3/4. We decided to align variation 12 to all the others by assigning to each inconstancy value the number of measures elapsed and not the offset in quarter lengths. This means that variation 12 is scaled on the time axis by a factor of 3, while all others by 2. This scaling

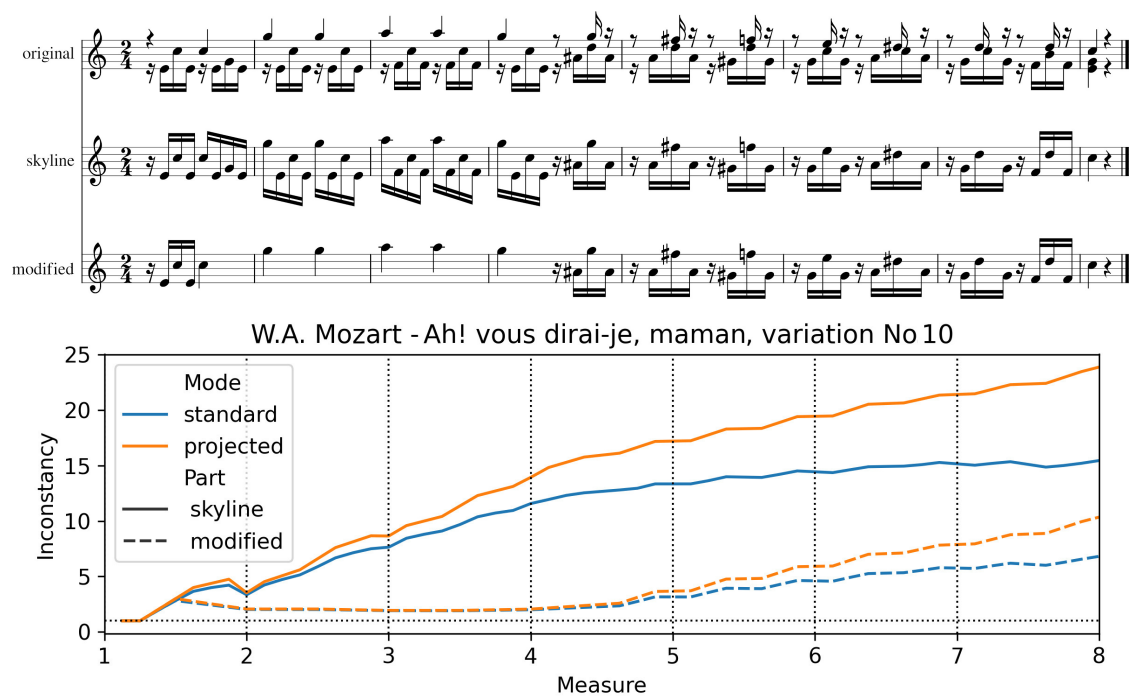


Figure 7. *Ah! vous dirai-je, maman*, variation No 10 by W.A. Mozart. The top staff is the original score for the right hand, the middle staff is the output of the skyline algorithm, and the last staff is from our modified version.

is performed *after* the calculation of the inconstancy. One could also decide to keep the values on the x -axis as they are or to stretch the music *before* calculating the inconstancy. Our choice was guided by the intuition that the metric structure offers a better case for alignment than the absolute duration of notes. However, we leave the investigation of this choice as an open research subject.

The inconstancy of the left hand is significantly and consistently larger than the one of the right hand. As a tentative explanation we notice that, in this particular set of variations, the right hand often has the melody, and the left hand the accompaniment. For example, take variations 7 and 10, the two for which the inconstancy for the right hand is the largest. A glance at the score⁸ shows that indeed, in these two cases, the right hand includes some typical accompaniment patterns similar to Alberti bass lines.

Therefore, the inconstancy could help in the detection of melodies versus accompaniment, which is an active research topic (see Chemillier 1987; Rizo et al. 2006; Madsen and Widmer 2007; and Guo, Herremans, and Magnusson 2019 for some approaches).

4. Applying the inconstancy to repertoire music

We use this section to give some ideas about how inconstancy could be used for more complex musicological tasks. In the first part, we make a statistical analysis of themes from the Barlow and Morgenstein corpus of musical themes; in the second part, we analyse entire pieces coming from common practice repertoire.

⁸Our reference edition is the one coming from the Werner Icking Music Collection and available at <https://imslp.org/wiki/Special:ReverseLookup/228371>.

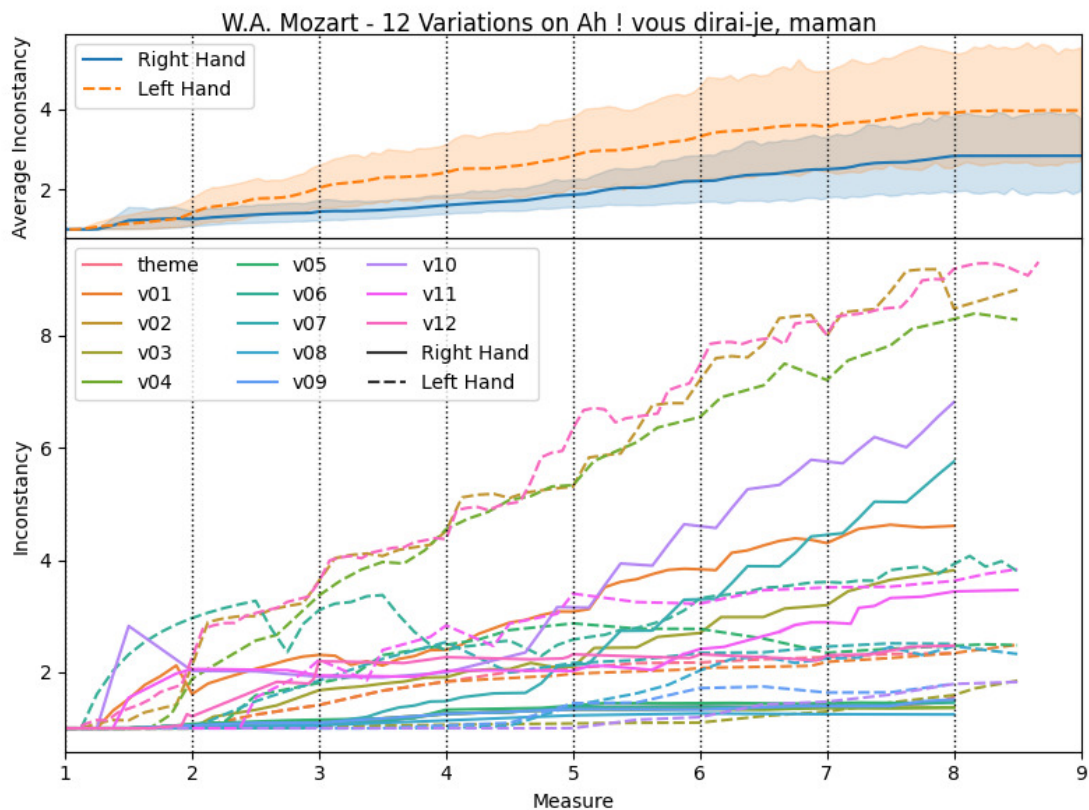


Figure 8. Inconstancy for the 12 Variations on *Ah! vous dirai-je, maman* by W.A. Mozart. The bottom plot shows the inconstancy of the theme and each variation separately, both for the right hand and the left hand. The top plot, instead, is the average over each hand of the interpolated inconstancies.

4.1. *An exploration of inconstancy on a large collection of themes*

When introducing a new feature, it is important to ask whether it helps in classification tasks. To address this question, we experimented with the different inconstancy functions proposed above on Barlow and Morgenstern's corpus. In order to obtain readable figures, we kept only the pieces of the composers that had more than 30 pieces in the corpus (this threshold value has been chosen after considering the distribution of the number of pieces for each composer in the corpus). It resulted in 65 composers from the Baroque Era to the XX^{th} century. Figure 9 shows the standard inconstancy value of the musical themes grouped by composer. This figure shows that the inconstancy values go mostly from 2 to 4. What strikes one, at first sight, is that some composers stand out from the rest, as Couperin and Scarlatti for instance, or Delibes to a lesser extent. A first look at the corresponding themes in the corpus indicates that it is due to the intensive use of ornamentation and arpeggios. Moreover, almost all the excerpts of Delibes in the corpus are from Ballets or dancing scenes in Operas. We did not further investigate these observations for the moment. Outliers are like some pieces of Kachaturian which are similar to arpeggios.

While observing Figure 9, we wondered whether the inconstancy could be in a certain way representative of the style of a composing period. We summarize our results in Table 1. We can see that the Baroque and Modern eras have a significantly higher average inconstancy than the other periods.




Figure 9. Standard inconstancy values grouped by composer in Barlow and Morgenstern’s themes corpus (outliers with an inconstancy value above 8 are not shown).

Period	Inconstancy	St.Dev	Average	#pieces
Baroque	3.193	0.040		1,068
Classical	2.437	0.015		2,720
Pre-Romantic	2.257	0.026		948
Romantic	2.419	0.016		3,247
Post-Romantic	2.343	0.030		683
Modern	2.611	0.025		1,497

Table 1. Summary of the average inconstancy over different periods. The standard deviation is also expressed not for the single piece but for the average over an entire period.

4.2. Sliding inconstancy for long music pieces

The results in Equation (5) seem to tell us that there is no point in studying the inconstancy of long pieces: if, after a certain duration, the inconstancy of music tends to a constant, it means that the new music we add does not provide any additional information. Therefore, the inconstancy of music is more informative when considered as a local feature.

In this respect, akin to spectrograms in audio content analysis, we study the value of inconstancy on sliding windows of music. We define a rectangular window of size W that moves along the time axis. Starting from the beginning of the piece, we take all notes with an offset between 0 and W and calculate their inconstancy. Then, we make the window hop rigidly to the right by a quantity H , so to frame all notes between H and $W + H$, and we repeat so until we reach the frame $[L - W, L]$, where L is the total length of the piece. The *inconstancy-o-gram* of music is thus the value obtained by this procedure written as a function of the frame index.

4.2.1. Identification of melodic vs. accompaniment parts

In Figure 10 we plot the distribution of the average value of the inconstancy on each frame for two corpora: the Barlow themes and Mozart’s string quartets, separated by

musical instrument. The results are in accordance with a few basic musicological observations. First of all, the violoncello parts tend to have the smallest inconstancy: that is compatible with the idea that cello parts tend to have longer notes, often repeated, to sustain harmonically the ensemble. On the other end of the spectrum, the first violin has the distribution that is the most similar to the one in the Barlow themes. Maybe because it has the most *melodic* parts in the quartet?

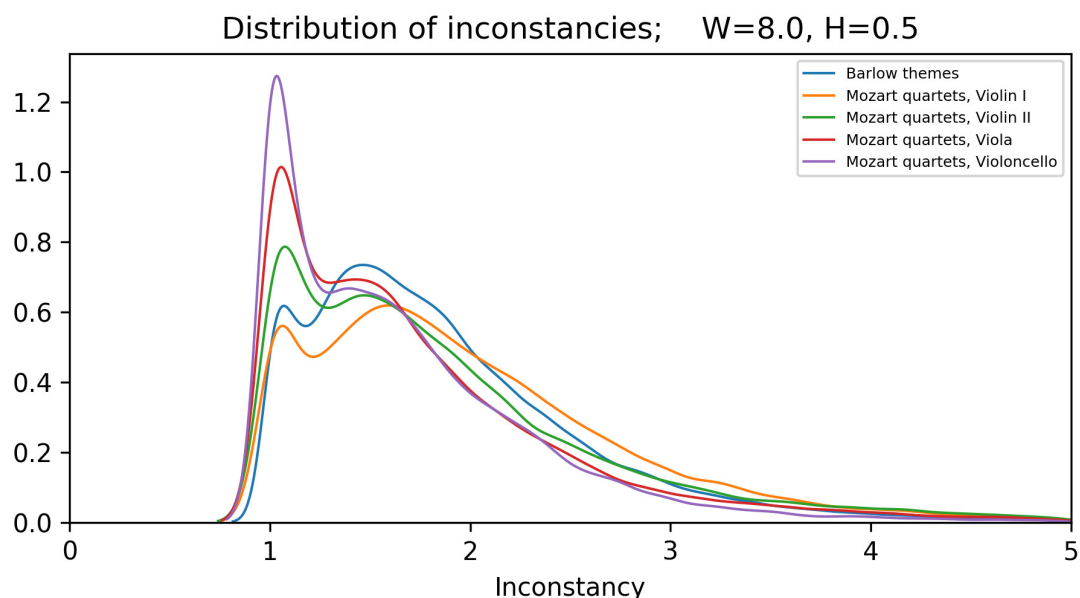


Figure 10. Distribution of the inconstancies on the Barlow themes corpus plotted against the distribution on Mozart quartets.

4.2.2. Structure analysis

In Figure 11 one can see the inconstancy-o-gram of the second movement of Mozart's K80 string quartet. The plotted inconstancy is the average over the inconstancies of the four parts (Violin I, Violin II, Viola, Violoncello). At the bottom of the figure, there is a line showing the structural segmentation of the music (a sonata form) and dots giving the position of perfect authentic cadences and root-position imperfect authentic cadences (Allegraud et al. 2019). It is apparent that the inconstancy is strongly correlated with the structural segmentation: for example, the behaviour of inconstancy during segments P to C and is very similar to the one between P' and C'. We conducted an experiment to introduce standard inconstancy as a feature for cadence detection in addition to the features of Bigo et al. (2018) for the same set of 21 Haydn quartets and saw indeed an improvement in the results: the number of false positives decreased from 28 to 21, showing that inconstancy is indeed related to cadential writing (maybe due to increased harmonic rhythm?). It would be interesting to study also the correlation between the inconstancy-o-gram of the four separate parts. This could highlight the presence of repeated patterns as peaks in the correlation function $\mathcal{S}_{\mathcal{I}_A \mathcal{I}_B}(t, t')$.

The inconstancy-o-gram can also be used to follow contrapuntal texture. For example, the inconstancy-o-gram applied to the Fugue in C major from J.S. Bach's Well-Tempered Clavier book I highlights the initial exposition of the subject in all four voices (figure not shown).

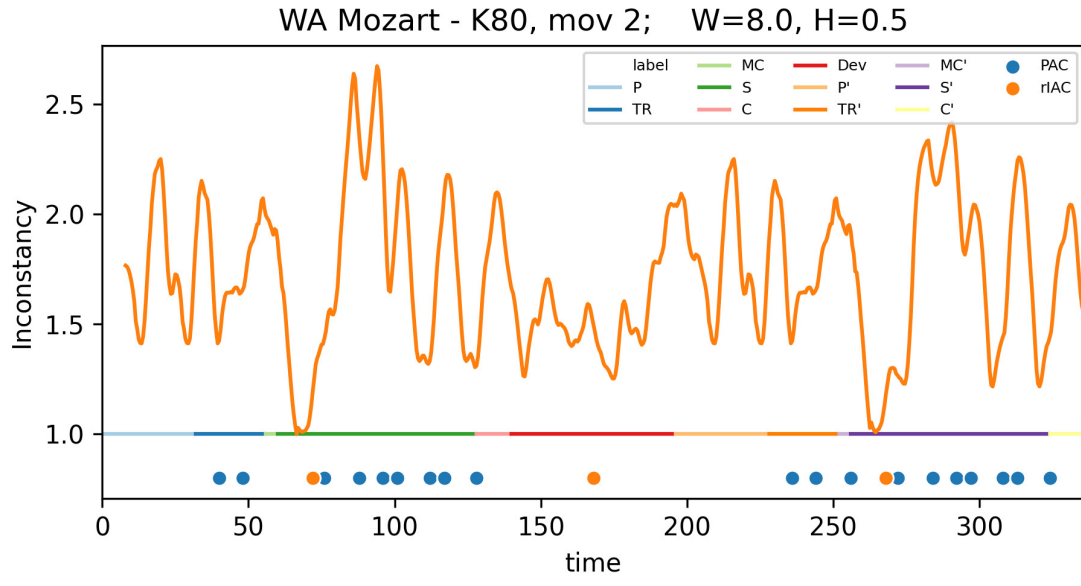


Figure 11. Inconstancy-o-gram of the second movement of the string quartet K80 by W.A. Mozart. The line has been smoothed with a Savitzky-Golay filter for better visualisation. The labels correspond to the several parts in a sonata-form: Primary (P) and Secondary (S) themes, Transition (TR), Conclusion (C), and Medial Caesura (MC), both in the exposition and in the recapitulation (same labels with primes), and Development (Dev). Labels PAC and rIAC correspond to structural breaks: perfect authentic cadence and root position imperfect authentic cadence.

5. Further work and conclusion

This paper introduced the notion of inconstancy for a monophonic musical line and characterized some of its mathematical properties.

The introduction of a new tool always opens the way to lots of possibilities for new investigations. Here we want to highlight the ones that seem more appealing to us. We begin with a mathematical question: is there a relationship between inconstancy and entropy, as was already discussed —although with a different name— in Denis and Crémoux (2002)? Which kind of information can be obtained from the entropy of a musical piece? Musically, it could also be worth studying how the inconstancy can be used to analyse and classify melodic contours, in comparison with other classifications such as the study of the melodic arch in (Huron 1996) on the Essen Folksong Collection (Schaffrath and Huron 1995). Preliminary experiments show that the inconstancy could be useful to identify specific contours that are not discriminated in Huron’s method, as highly oscillating melodies, which could therefore guide us to an improvement in the classification of melodies, at a small expense in the complexity of the system. Another task that seems particularly interesting to us concerns the detection of phrase boundaries, and the use of the inconstancy as an input feature to improve the accuracy of the LBDM model (Cambouropoulos 2001). The analysis of structure is still one of the key challenges in computational musicology, and we believe that the inconstancy can prove useful not only to detect boundaries but also to classify the musical content inside the boundaries, as we have briefly seen from the experiment on cadence detection reported in Figure 11. More generally, we think that inconstancy could help to detect changes of texture, allowing to capture the impression of stability or movement of the music. The inconstancy could also play a role in machine learning models for music generation and analysis as a constraint that generated melody should statistically satisfy. And, finally, a practical question is

how to apply inconstancy to polyphonic music: by separately computing inconstancies of, say, principal melody and bass line(s) if relevant? by “averaging” different voices? by dismissing the notes with a long duration and only keeping the non-overlapping shortest ones? A definitive answer to this question would render this new tool applicable to all music pieces.

Acknowledgements

Author Jean-Paul Allouche wants to thank Carlos Agon and Moreno Andreatta for interesting discussions on the application of inconstancy to music analysis. Author Gianluca Micchi thanks Guillaume David for his help with the statistical tests of significance. We thank the editor and the anonymous reviewers, as well as Mathieu Giraud, for their insightful comments.

Funding

This work is partially funded by French CPER MAuVE (Région Hauts-de-France) and by a grant from the French Research Agency (ANR-11-EQPX-0023 IRDIVE).

References

- Allegraud, Pierre, Louis Bigo, Laurent Feisthauer, Mathieu Giraud, Richard Groult, Emmanuel Leguy, and Florence Levé. 2019. “Learning Sonata Form Structure on Mozart’s String Quartets.” *Transactions of the International Society for Music Information Retrieval* 2 (1): 82–96.
- Allouche, Jean-Paul, and Laurence Maillard-Teyssier. 2011. “Inconstancy of Finite and Infinite Sequences.” *Theoretical Computer Science* 412 (22): 2268–2281. <http://www.sciencedirect.com/science/article/pii/S0304397511000181>.
- Barlow, Harold, and Sam Morgenstern. 1948. *A Dictionary of Musical Themes*. Crown Publishers.
- Bigo, Louis, Laurent Feisthauer, Mathieu Giraud, and Florence Levé. 2018. “Relevance of Musical Features for Cadence Detection.” In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR 2018)*, <https://hal.archives-ouvertes.fr/hal-01801060>.
- Buccoli, Michele, Massimiliano Zanoni, Augusto Sarti, Stefano Tubaro, and Davide Andreoletti. 2016. “Unsupervised Feature Learning for Music Structural Analysis.” In *2016 24th European Signal Processing Conference (EUSIPCO)*, Aug., 993–997. ZSCC: 0000003 ISSN: 2076-1465.
- Cambouropoulos, Emilios. 2001. “The Local Boundary Detection Model (LBDM) and its Application in the Study of Expressive Timing.” In *Proceedings of the International Computer Music Conference (ICMC’2001), La Habana, Cuba edited by A. Schloss and R. Dannenberg, (International Computer Music Association, San Francisco)*, 290–293. <https://ofai.at/papers/oefai-tr-2001-11.pdf>.
- Carsault, Tristan, Jérôme Nika, and Philippe Esling. 2018. “Using Musical Relationships between Chord Labels in Automatic Chord Extraction Tasks.” In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR 2018)*, Paris, France, Sep., 18–25. <https://hal.archives-ouvertes.fr/hal-01875784>.
- Chemillier, Marc. 1987. “Monoïde libre et musique : deuxième partie.” *RAIRO - Theoretical Informatics and Applications - Informatique Théorique et Applications* 21 (4): 379–417.
- Chen, Tsung-Ping, and Li Su. 2018. “Functional Harmony Recognition of Symbolic Music Data with Multi-task Recurrent Neural Networks.” In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR 2018)*, 90–97.
- Chen, Tsung-Ping, and Li Su. 2019. “Harmony Transformer: Incorporating Chord Segmentation into Harmony Recognition.” In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR 2018)*, 259–267.
- Denis, Alain, and Francis Crémoux. 2002. “Using the Entropy of Curves to Segment a Time or Spatial Series.” *Mathematical Geology* 34 (8): 899–914. <https://doi.org/10.1023/A:1021302922108>.

- Guo, Rui, Dorien Herremans, and Thor Magnusson. 2019. “Midi Miner: A Python library for tonal tension and track classification.” *ArXiv* abs/1910.02049.
- Huron, David. 1996. “The melodic arch in Western folksongs.” *Computing in Musicology* 10: 3–23.
- Huron, David. 2013. “On the Virtuous and the Vexatious in an Age of Big Data.” *Music Perception* 31 (1): 4–9. <https://doi.org/10.1525/mp.2013.31.1.4>.
- Ju, Yaolong, Nathaniel Condit-Schultz, Claire Arthur, and Ichiro Fujinaga. 2017. “Non-Chord Tone Identification Using Deep Neural Networks.” In *Proceedings of the 4th International Workshop on Digital Libraries for Musicology*, DLFM ’17, New York, NY, USA, 13–16. Association for Computing Machinery. <https://doi.org/10.1145/3144749.3144753>.
- Kirlin, Phillip B., and Jason Yust. 2016. “Analysis of analysis: Using machine learning to evaluate the importance of music parameters for Schenkerian analysis.” *Journal of Mathematics and Music* 10 (2): 127–148. <https://doi.org/10.1080/17459737.2016.1209588>.
- Madsen, Søren Tjagvad, and Gerhard Widmer. 2007. “Towards a Computational Model of Melody Identification in Polyphonic Music.” In *Proceedings of IJCAI’07*, San Francisco, CA, USA, 459–464. Morgan Kaufmann Publishers Inc.
- McKay, Cory, Julie Cumming, and Ichiro Fujinaga. 2018. “JSymbolic 2.2: Extracting Features from Symbolic Music for use in Musicological and MIR Research.” In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR 2018)*, 348–354.
- Micchi, Gianluca, Mark Gotham, and Mathieu Giraud. 2020. “Not All Roads Lead to Rome: Pitch Representation and Model Architecture for Automatic Harmonic Analysis.” *Transactions of the International Society for Music Information Retrieval* 3 (1): 42–54.
- Mishra, Saumitra, Bob L Sturm, and Simon Dixon. 2017. “Local Interpretable Model-agnostic Explanations for Music Content Analysis.” In *Proceedings of the International Society for Music Information Retrieval Conference 2017*, 537–543.
- Müllensiefen, Daniel, and Geraint Wiggins. 2011. “Polynomial functions as a representation of melodic phrase contour.” In *Systematic Musicology: Empirical and Theoretical Studies*, 63–88. Peter Lang.
- Nápoles López, Néstor, Laurent Feisthauer, Florence Levé, and Ichiro Fujinaga. 2020. “On Local Keys, Modulations, and Tonicizations: A Dataset and Methodology for Evaluating Changes of Key.” In *Proceedings of the 7th International Conference on Digital Libraries for Musicology*, DLFM 2020, New York, NY, USA, 18–26. Association for Computing Machinery. <https://doi.org/10.1145/3424911.3425515>.
- Pauwels, Johan, Ken O’Hanlon, Emilia Gómez, and Mark B. Sandler. 2019. “20 Years of automatic chord recognition from audio.” In *Proceedings of the International Society for Music Information Retrieval (ISMIR)*, Delft, Netherlands, 04/11/2019, 54–63. <http://hdl.handle.net/10230/42773>.
- Rizo, David, Pedro J. Ponce De León, Antonio Pertusa, and Jose M. Iñesta. 2006. “Melody track identification in music symbolic files.” In *Proceedings of the 19th International FLAIRS Conference*, 254–259.
- Rydén, Jesper. 2020. “On features of fugue subjects. A comparison of J.S. Bach and later composers.” *Journal of Mathematics and Music* 14 (1): 1–20. <https://doi.org/10.1080/17459737.2019.1610193>.
- Schaffrath, Helmut, and David Huron. 1995. *The Essen folksong collection in the humdrum kern format*. Menlo Park, CA: Center for Computer Assisted Research in the Humanities.
- Tymoczko, Dmitri. 2011. *A Geometry of Music: Harmony and Counterpoint in the Extended Common Practice*. New York: Oxford University Press.
- Uitdenbogerd, Alexandra, and Justin Zobel. 1999. “Melodic matching techniques for large music databases.” In *Proceedings of the seventh ACM international conference on Multimedia (Part 1)*, MULTIMEDIA ’99, Orlando, Florida, USA, Oct., 57–66. <https://doi.org/10.1145/319463.319470>.
- Ullrich, Karen, Jan Schlüter, and Thomas Grill. 2014. “Boundary Detection in Music Structure Analysis using Convolutional Neural Networks.” In *Proceedings of the International Society for Music Information Retrieval Conference 2014*, 417–422.