



HAL
open science

Automating Motivic Analysis through the Application of Perceptual Rules.

Olivier Lartillot, Emmanuel Saint-James

► **To cite this version:**

Olivier Lartillot, Emmanuel Saint-James. Automating Motivic Analysis through the Application of Perceptual Rules.. MIT Press; Walter B. Hewlett, Eleanor Selfridge-Field. Music Query: Methods, Models, and User Studies, 13, pp.73-92, 2004, Computing in Musicology, 0-262-58256-2. hal-01161136

HAL Id: hal-01161136

<https://hal.science/hal-01161136v1>

Submitted on 8 Jun 2015

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.

Automating Motivic Analysis through the Application of Perceptual Rules

Olivier Lartillot

Musical Representations Team
Ircam – Centre Pompidou
Place Igor-Stravinsky
75004 Paris, France
Olivier.Lartillot@ircam.fr

Emmanuel Saint-James

LIP6/SRC
Université Pierre et Marie Curie
4, place Jussieu
75255 Paris, France
Emmanuel.Saint-James@lip6.fr

Abstract

The musical discourse may be described in terms of an intricate flow of local groupings. Such groupings, whose perception does not always reach a state of explicit awareness, mostly remain in an informal form, except the most predominant of them, which contribute to more global constructions and will be remembered as the characteristic thematic materials of the musical piece. There have been some attempts, in particular Reti's thematic analysis, to explicitly describe music at this level of detail. But such non-reductionist approach of music analysis, facing huge complexity, desperately needs automation and objectivity.

Current researches in Musical Pattern Discovery, which may be considered as first steps towards this ideal, hardly discover the basic musical structures expected by musicologists. This failure stems from the fact that current formalizations of musical patterns do not take plainly into account the essential characteristics of music as a perceptual phenomenon.

We propose a new approach of Musical Pattern Discovery, founded on perceptual heuristics, with the ideal aim of making explicit all the structural details that we more or less implicitly perceive. Basic principles and algorithms are described and illustrated, and first promising results are shown.

1 Perceiving the Motivic Dimension of Music

1.1 Toward Motivic Analysis

One main reason for the growing interest of musicologists in computer tools is that large databases can be automatically processed. Huge corpuses can be searched for musical entities, either through a matching of a specific query or through an automated discovery of new configurations. Such user-defined queries and computer-defined discoveries are constrained to belong to certain musical categories, depending on the specific purposes for which these computer applications are designed.

For instance, such inquiry may focus on the *melody*, which is considered as one of the most salient aspects of music, with the view to effectuate a discriminating characterization of works (see Eleanor Selfridge-Field 1997-1998). Multiple avatars of a particular melody may be retrieved among the corpuses, be they exactly similar, slightly distorted, or significantly developed. It has been acknowledged, however, that the concept of melody cannot be applied universally, since some musical genres content themselves with constantly proliferating basic gestures instead [Figure 3].

Another musical concept, onto which may these computer-aided analyses be centered, is the *theme*. Despite its confusing semantic similarities with the concept of melody, the thematic dimension of music may be in fact considered as an orthogonal vision. A theme is indeed any entity that characterizes a work because of its reiterated presence. A theme may be melodic, but also rhythmic, harmonic, etc. Reversely, a melody that does not share any specific redundancy with the rest of the work may be considered as non-thematic. What is interesting here is that the concept of theme is, contrary to that of melody, less dependent on cultural knowledge.

Since a theme is a specific configuration that is repeated several times throughout the piece, one may deepen the analysis by focusing on such basic configurations themselves, be they thematic or not. Such approach, initiated by Rudolph Reti (1951), and considering music analysis as a discovery of the specificity of each local aspect of the piece, contrasts with the reductionist demarche that is usually undertaken by traditional musicology. Reti's thematic analysis has been much criticized, maybe because it desperately needs the help of automation. Indeed, the detailed analysis of local aspects easily "degenerates into a purely mechanical exercise in which the score is analyzed without ever really being read properly" (Nicholas Cook 1987). Moreover, the foundation of his methodology, revealing his personal conception of music aesthetics, lacks objectivity. A computer implementation of such approach would require an explicit description of these mechanisms.

1.2 The Unattainable Deletion of the Listener

Such analysis on the thematic level of music suggests another interest of musicology in computer: as a tool for explicating all the pertinent structures that may be discovered in the score. Unfortunately, such task is far from easy to achieve pertinently. Today, only human intelligence is able to catch more or less the subtle configurations in music. Does such idealistic artificial musical intelligence make sense, or should it be considered as a utopian dream that cannot be practically envisaged? This requires a precise explanation of the task of objective characterization of musical structures, and questions on the possibility of exhaustive discovery of such entities.

This problem may be considered from a purely structural point of view. Jean-Jacques Nattiez's "*analysis at the neutral level*" (1990) focuses on the "immanent configurational properties of a musical work", which is supposed to exist independently of the "*poietic level*" of compositional procedures and intentions, and the "*aesthetic level*" of perceptual processes. According to such approach, musical entities may *exist* primarily in the score even before the listening or the analyzing process discovers them. However, such distinction between subjective perception and objective analysis appears less evident when going into details of what any neutral level analysis really consists of.

The emblematic example of neutral level analysis is the *paradigmatic analysis*, initiated by Nicolas Ruwet (1987), and based on a multi-levelled research of repeated patterns. Such approach, if we temporally suppose restricted to exact repetitions, may appear entirely objective at first sight. But where does the simple definition of a pattern as a contiguous sequence of notes, and not as a set of notes that are scattered throughout the score, come from, if not from cognitive properties?¹ For instance, an exact repetition of a pattern may equivalently be considered as an overlapping of couples of notes [Figure 1]. This second dual representation is implicitly discarded simply because it does not correspond to what we actually hear.



Figure 1. Mathematically speaking, a succession of two identical patterns of contiguous notes (solid lines) is equivalent to its dual representation as an overlapping of several identical patterns of two notes (dashed lines).

Things get worse when approximate repetitions have to be considered. As explained Emiliós Cambouropoulos, "if similarity (i.e. not merely exact

¹ According to Ruwet himself, paradigmatic analysis *does* take into account perceptual principles.

repetition) is taken into account then analysis at the neutral level becomes unwieldy because any two musical sequences are similar in some respect. Analysis at the neutral level is useful only if guided by some sort of heuristics – for instance, based on general cognitive principles.” (1998:12)

Such conclusions, calling into question the ideal of “neutrality”, which, according to Nattiez, “means both that the poietic and aesthetic dimensions of the object have been ‘neutralised’, and that one proceeds to the end of a given procedure regardless of the results obtained” (1990:13), would rather lead us to defend the idea that musical structures cannot be separated from compositional or listening strategies.

1.3 Explicating the Implicit

There remain two levels along which the numerous musical structures may be envisaged: either poietic or aesthetic. Reti, since he was himself a composer, envisions thematic structure mainly on the first level: “the thematic phenomena are so manifold and complex that in a sense they evade academic tabulation. Though they can perhaps be described, they can hardly be comprised in an actual “system.” They are too intimately connected with the creative process itself.” (1951:233) But evidently, from the listener’s point of view, composer intentions may show through the score.

Therefore, what is the use of musical analysis if a mere perception of music would suffice? The superiority of analysis over perception may be explained by a frankly elitist argument, stating that musical analysis shows the way we have to perceive music. Or it may be explained by the idea that analysis explicates what we perceive in a more or less explicit way. Once these inferences are made explicit, their perception will be alleviated and more enjoyed. But in the same time, the potential of music perception is so rich that it seems really difficult, at first glance, to reduce it into a collection of explicit inferences.

2 Basic Principles of Pattern Perception

This project of automated music analysis may appear particularly tempting to musicology. However, for all the progress in computer science and artificial intelligence, this remains a Utopia. The mere quest of pertinent musical patterns, which are the basic blocks of musical structures, remains an intricate challenge that has not been met yet.

2.1 Characterizations of Patterns

A musical pattern may be defined in three different ways: first its existence may be explicitly and logically deduced from the constructive rules of a particular style. For instance, works of the classical style may roughly be considered as a hierarchical construction of patterns (such as antecedent and consequent ones) whose length is expressed in a fixed number of bars. However, even when actual musical works deviate from this emblematic model, their musical patterns are still understood by the listener. That is why the general characterizations of patterns may be considered as independent of style.

Pattern may be discovered in a purely inductive way, following perceptual processes. Results offered by psychological studies such as Gestalt theory have been broadly applied in musicology, by, among others, Meyer (1956) and Narmour (1990). In particular, groupings may result from segmentations that are ruled by local properties of the musical discourse. For instance, segmentation may be introduced between entities that contrast one with the other along their pitch, time onset, duration or intensity dimensions (see Lerdahl and Jackendoff, 1983, Temperley, 2001 and Cambouropoulos 1998).

Alternatively, a musical pattern may be defined as a set of notes that is repeated several times throughout the score. This last criterion, which has been largely studied in linguistics (where paradigmatic analysis may be considered as a musical application, see paragraph 1.2), conflicts with the second criterion of local segmentation. In particular, a musical pattern may be implicitly built through contrastive aggregation. For instance [Figure 2], the initial leap between the first two notes of a pattern, although triggering a contrastive idea [dashed lined], may be the significant element that characterizes the beginning of the pattern itself [solid lines]. A pattern is therefore not a conservation of sameness, but rather a travel through contrasts.



Figure 2. A pattern may feature contrastive steps.

In the remaining of this study, we will mainly focus on this last repetition-based criterion², but Gestalt principles will be considered too. All supplementary

² Lerdahl & Jackendoff's generative theory includes the concept of repetition, under the term *parallelism*. They agree that "a set of preference rules for parallelism has to be developed." They recognize however "not to be prepared to go further" and "feel that their incapacity to flesh out this notion of parallelism is an important failure in this attempt to formulate a fully explicit theory of musical understanding" (1983:53). Similarly: "Because motivic analysis is influenced by a variety of factors, it seems likely that it, too, would be

knowledge, such as harmony or meter, will be discarded. Indeed, repeated patterns may appear on harmonically and metrically unstable places, like overlapping theme in a fug stretto.

2.2 Cognitive Descriptions of Patterns

The researches dedicated to such aspect of musical pattern, constituting a discipline called Musical Pattern Discovery, mostly agree to a same unique basic methodology for extracting patterns, that consists in a detection of similarity of succession of musical elements (either notes or intervals) (see for instance Tim Crawford et al. 1997-1998). Such criterion, although a necessary condition for pattern determination, is however not sufficient to insure a sound characterization. Simple examples may easily show that similarities of succession of notes do not necessary characterize pertinent patterns. In Bach's *Prelude in C*, BWV 846 [Figure 3], some repetitions of successions of 8 notes may correspond to actual motives (such as the solid lined ones), both some may correspond to non-pertinent ones too (dashed lines).



Figure 3. Two exact repetitions, one corresponding to pertinent patterns (solid lines), the other to non-pertinent ones (dashed lines).

If previous example was not convincing enough, this simpler example [Figure 4], which only features exact repetitions, may induce the inference of numerous non-pertinent patterns, due to the presence of a trill.



Figure 4. Non-pertinent patterns induced by the presence of a trill.

In order to avoid such bad inferences, Cambouropoulos (1998), who considered such difficulties, proposes to add a specific constraint, stating that overlapping of patterns should be avoided. However, such constraint, specified upon the general classes of pattern directly and not on the specific occurrences, cannot avoid the inference of patterns like FGF [Figure 4] or the dashed lined patterns [Figure 3].

Such difficulty may be solved by considering the fact that patterns stem from conceptual inferences that are progressively processed during the incremental

amenable to a preference rule approach, this [is not] attempted” by Temperley’s approach (2001:335).

listening of the piece. Each time a new note is considered, the set of inferences that are currently in process constitutes a context, which induces constraints upon the candidates for new inferences. In this way, non-pertinent inferences described in previous examples may be avoided.

2.3 Formalization of Patterns: Classes and Occurrences

As we need to design a system that incrementally scans the score and that progressively produces inferences that should depend on the score and on what has been already inferred, it may be convenient to construct these inferences directly on the score. First of all, we have to clearly distinguish the concepts of pattern class and pattern occurrence. A *pattern class* is an abstract concept that unifies all patterns of the score that are considered as similar all together. These patterns will be called *pattern occurrences* of the pattern class.

The characteristics of a pattern class are expressed through a list of properties shared by notes of the patterns and by the intervals between successive notes. Moreover, if we restrict our scope to monophonic music in a first approach, the notes of the pattern are totally ordered by chronology. Therefore, the entity of a pattern, be it a pattern class or a pattern occurrence, may naturally be formalized through a chain of states, where successive states corresponds to a successive notes of the pattern, and where each transition between two successive states corresponds to an interval between two successive notes [Figure 5].

Concerning the *pattern class chain* (PCC), states and transitions feature the characteristic properties shared by the respective notes and intervals of its corresponding pattern occurrences. The *pattern occurrence chain* (POC) interfaces the specific notes in the score with the associated PCC: each state of the POC is linked to one note of the score and in the same time to one state of the PCC.




Figure 5. The two POCs (black circles) interface notes in the score with the corresponding states in the PCC (white circles).

Such formalization through graph theory enables to easily implement the rules that should be added to the framework, in order to insure pertinent inferences. For instance, there should exist a rule, generalizing our previous remarks concerning overlapping (see paragraph 2.2), and stating that suffixes of patterns cannot be considered as new patterns, since they are already subsumed by current patterns [Figure 6].




Figure 6. Suffixes (bottom POCs and PCC) cannot be considered as new pattern classes, since the associated notes of each of its POCs already belongs to the POCs of the whole pattern (top POCs and PCC).

2.4 Patterns Associations

Music features multiple levels of pattern descriptions. Notes may indeed belong to several possible patterns in parallel. When discovering another occurrence of a pattern class that was associated to another pattern class, we may expect to retrieve the same association between these pattern classes. Therefore, pattern association may induce pattern *expectation*, as would say Leonard Meyer (1956).

For this purpose, the association between patterns may be directly represented on the PCCs themselves. Every time a note of a pattern occurrence is associated to another pattern occurrence, on the corresponding state in each PCC is associated a new POC associated to the other PCC. Such pattern association discovery induces a pattern association expectation rule, stating that every time a new occurrence of such pattern class is discovered, possible associated pattern classes are also expected. Such expectation may be formalized through the instantiation of new hypothetical POCs whose only first state is represented, waiting for further extension.

For instance, a pattern may include another sub-pattern. The 8-note pattern of Bach's *Prelude* in *C* features a repetition of two 3-note patterns [Figure 7]. Therefore, to the 8-note pattern class itself is added two 3-note pattern occurrences.




Figure 7. The 8-note pattern of Bach’s *Prelude* in *C* includes a repetition of two 3-note patterns.

2.5 Pattern Reiteration

Music usually features successive repetitions of a same pattern class [Figure 3]. If no new mechanisms were added, the system would consider each possible concatenation of these successive patterns as a new pattern. As told previously (see Paragraph 2.2), these inferences, not corresponding to human judgments and leading to combinatorial explosion, should be forbidden.

It may be remarked that such pattern repetition is a special case of pattern association. If each pattern is extended with the first note of the succeeding pattern, then this last note of such extended pattern may be associated to the first note of the same pattern class [Figure 8]. This means that, in the extended pattern class, the last state is linked to the first state. The idea of pattern cycling is therefore explicitly represented.

The first note of each new occurrence, as soon as it appears, is immediately associated to a new pattern occurrence chain. An additional mechanism prevents any pattern, whose first note is also the last note of another occurrence of the same pattern, to be extended further.




Figure 8. The last note of the 9-note pattern is linked to its first one.

2.6 Meta-Pattern of Patterns

The succession of pattern occurrences may be examined in exactly the same way that the previously considered succession of simple notes, provided that these pattern occurrences belong to the same pattern class³. For this purpose, a concept of “interval” between two successive pattern occurrences, that generalizes the traditional notion of interval between two successive notes, is defined. Such generalized interval simply consists of the list of all interval relationships between the two respective notes of *same rank* in each pattern occurrence.

³ In future works, such a strong restriction may be replaced by a weaker one: for instance, pattern prefixes may be accepted too.



Figure 9. In a generalized interval between pattern occurrences, notes of same rank are compared.

Meta-pattern of pattern occurrences will then be inferred when the successive generalized intervals of several successions of pattern occurrences are similar. It has to be remarked, here, that the similarity between two generalized intervals may consist of a partial similarity along one or several ranks of the associated pattern occurrences only. In this way, melody formed by the first note of each successive pattern occurrence, for instance, is considered (as in bars 8-11 and 15-18 of the Bach's *Prelude*). If the pattern occurrences are heard as chord arpeggiations, this generalized vision of pattern discovery may simply be understood as a horizontal approach of the successive pattern occurrences, whereas the previous vision was the vertical one.

Exact repetition of pattern occurrences may be considered as generalized "unison". The *Prelude* may then be modeled as a "perpetual" reiteration of a meta-pattern class that consists of two 8-note patterns in generalized unison [Figure 10]. Moreover, this meta generalization may be operated recursively and meta-pattern of meta-pattern may be considered too! Similar chord successions in the *Prelude* are special occurrences of such meta-patterns [Figure 11].




Figure 10. Each bar of the Bach's *Prelude* contains a same generalized interval of "unison".




Figure 11. The *Prelude* also features repeated meta-patterns of meta-patterns of 8-note patterns.

3 Modeling the Emergence of Similarity

3.1 Towards an Inductive Inference of Identities

So far, we have only considered exact repetition of patterns. The inference of identity relationships between patterns that are only similar is not obvious at all. One common idea consists of looking for identity along different “*viewpoints*” (see Darrell Conklin et al. 1995), or “level of abstraction” (Dannenberg 2002) of the score.

Simply transposed patterns [Figure 12] may be detected by considering each pattern in its own pitch reference. For example, if patterns are described not with absolute pitch, but with relative pitch whose reference is the absolute pitch of the first note of the pattern, then such descriptions of transposed patterns are exactly identical.



Figure 12. A pattern and one transposition of it.

As for transpositions, slower and faster patterns [Figure 13] may be considered as identical one to the other if a relative temporal representation is

considered. For this purpose, the quotient between the duration of current note and the duration of the first note is preferred to the absolute duration of current note.



Figure 13. A pattern and one time stretching of it.

But real music features much more complex transformations. In particular, pitch and temporal distortions may appear locally inside patterns. In order to handle such plasticity, more relative viewpoints of the pattern are considered. One commonly used description is the contour, which consist of the direction of the interval between successive notes (downward, upward, or a repeat). This representation is so loose that non-pertinent repetitions may also be detected. Patterns with same contour (for instance, up, down, down, down) [Figure 14] may be considered as nearly identical [first and second pattern], or as significantly different [first and third].



Figure 14. A pattern, a slightly distorted version of it, and a different pattern with same contour description.

In fact, locally distorted patterns do not generally feature any explicit repetition. Repetitions are only hypothetical characteristics that have to be actively *induced* during perception (see Lartillot 2002). What have also to be taken into consideration are the local characteristics such as the quantitative distance between different intervals, or couple of successive notes. In previous example [Figure 14], the first two patterns are similar because the distances between successive intervals are similar.

3.2 Local Contexts

If a pattern induction algorithm has to mimic human capabilities, their basic principles need first to be described. In particular, we need to understand how a human listener, once beginning to hear the second occurrence of a pattern, is able to suddenly remember that what is being played has already appeared previously — exactly or similarly —, even when such pattern was not already explicitly defined.

Such cognitive capabilities seem to rely on the general characteristics of associative memory, where knowledge is indexed by their content instead of being referenced according to any arbitrarily defined memory address. As pattern may be recalled even before being explicitly discovered, there exists a

reproductive memory that associates local succession of notes that are similar one to the other. Patterns may be defined as successions of local similarities.

These local similarities are progressively built from single intervals. Distances are computed first between single intervals, then between succession of intervals, or patterns.

Let n_1, n_2, n_3, n_4 be four notes whose respective pitches are p_1, p_2, p_3, p_4 and respective durations d_1, d_2, d_3, d_4 . We propose, in a first approach, to formalize the perceptual distance between two intervals (n_1, n_2) and (n_3, n_4) as a weighted product of a pitch distance and a duration distance:

$$D((n_1, n_2), (n_3, n_4)) = (\text{abs} [(p_2 - p_1) - (p_4 - p_3)] + 1) \square (\max [d_1/d_3, d_3/d_1])^{0.7}$$

Here, only the duration of the first note of each interval is taken into consideration, since this duration is also the temporal distance between the successive notes. We may also remark that in music, pitches are subtracted, whereas durations are divided. In future works, we would like to replace such *ad hoc* measure with more firmly grounded modeling.

3.3 Step 1: Discovering Similar Contexts

Every successive local interval has to be memorized in an associative memory that is able to retrieve any interval similar to a query. For this purpose, a hash-table associates for each interval parameter the set of its occurrences in the part of the score that has already been analyzed. For any current local interval i_1 , each similar old local interval i'_1 is retrieved through a lookup of the hash-table.

If similarity between i_1 and i'_1 exceed a certain threshold, previous local interval i_0 that precedes i_1 is considered, and compared to previous local interval i'_0 that precedes i'_1 [Figure 15]. If the similarity between sequences (i_0, i_1) and (i'_0, i'_1) exceeds another threshold, then a pattern class, that consists of this succession of two intervals, is inferred. Such approach may then be generalized to sequences longer than 2 intervals.




Figure 15. During step #1, when two intervals i_1 and i'_1 are similar, preceding intervals i_0 and i'_0 are also compared.

Here an additional constraint for the choice of similar context should be introduced: namely, that old sequences have first to be pre-segmented by Gestalt rules, such as a contour discontinuity. Indeed, if the old sequence (i'_0, i'_1) , even if strictly identical to (i_0, i_1) , is in fact included in continuously upward contours [Figure 16, first dashed lines on bar 2], then it seems that it cannot be extracted from its context. It can only when such sequence features a contour discontinuity [first solid lines on bar 2], even when the resulted sequence is much less similar.



Figure 16. Only pre-segmented sequences may be retrieved.

There should also be a distinction between two types of similarity relationships: current sequence may be associated either to most recent — but not particularly similar — sequences, or to more ancient and very similar sequences. First type of relationship enables the constitution of successive pattern occurrences of same class and hence meta-patterns (see paragraph 2.6), whereas second type of relationship accounts for the discovery of firmly constituted patterns that are repeated more infrequently throughout the score, and in particular the exact repetition of parts of the piece, such as a *da capo*. Close patterns can be dissimilar, since they are easily accessed through short-term memory, whereas distant patterns have to be very similar, in order to be retrieved amongst the numerous items of long-term memory.

3.4 Step 2: Prolonging Contexts

Once the similarity between the beginnings of two patterns (i'_0, i'_1) and (i_0, i_1) have been detected, the inference of the similarity between the continuations of each pattern i'_2 and i_2 [Figure 17] is far easier to be processed, since there is no need to discover new contexts any more. As present context induces *expectation*, in Meyer sense (1956), the only necessary conditions for i_2 to be effectively related to i'_2 is simply that they share same contour. In our approach,

therefore, contour is effectively considered as a pertinent parameter for pattern discovery. But thanks to the additional restriction that pattern initiation is ruled by a more careful and less tolerant heuristic, which takes into account interval distance instead of simple contours and Gestalt segmentation (see paragraph 3.3), a large amount of false positive similarities are avoided.




Figure 17. Step #2 simply consists of a comparison of each successive note of one pattern with the respective note of the other pattern.

Each successive note of the currently heard repetition is tentatively related with a possible successive note of the previously heard occurrence [Figure 17]. In a pattern, each note may be disposed relatively to the position of its preceding note, or also relatively to the position of the first note of the pattern, and sometimes even relatively to the position of another particular previous note of the pattern. In the 8-note pattern of Bach's *Prelude* in *C* [Figure 18], the similarity between the third note of each pattern is explained by its relative position with respect to the first note, and not the second note, since the two intervals between the second and the third note of each pattern are particularly different, whereas the two intervals between the first and the third note of each pattern are less dissimilar. On the contrary, the similarity between the fourth and fifth notes of each pattern may simply be explained by their position relatively to their preceding notes.



Figure 18. Local referential relationships between pattern notes.

That is why the position of each note in a pattern may be considered relatively to each possible previous note in the pattern, in order to find the minimum dissimilarity. More precisely, the distance between the two successions of notes (n_1, n_2, \dots, n_L) and $(n'_1, n'_2, \dots, n'_L)$ is equal to:

$$D((n_1, n_2, \dots, n_L), (n'_1, n'_2, \dots, n'_L)) = \min_{1 \leq i < L} D((n_i, n_L), (n'_i, n'_L))$$

3.5 Discovering Further Occurrences

Once a pattern class has been discovered, its further occurrences should all be subsumed under the same pattern class. Therefore the discovery of these further occurrences cannot obey strictly to the pattern discovery algorithm described by previous steps #1 and #2. However, the distinction between context discovery and context prolonging still prevails.

In step #1, when a similarity has been discovered between two different contexts, and before deciding to create any new pattern class, we have to make sure that such context does not already exist in the beginning of one of the set of all discovered pattern classes (or more simply to the set of pattern classes associated to the past context). If there does exist such pattern class, a new pattern occurrence simply associates the new discovered context with the retrieved pattern class.

Step #2 of pattern occurrence discovery significantly differs from step #2 of pattern class discovery. Since the beginning of currently discovered pattern occurrence is already associated to a pattern class, each of its successive candidate continuations may simply be compared to the successive continuations along the pattern class. In this case, current pattern occurrence does not need to be compared to old occurrences.⁴

4 *OMkanthus*, an *OpenMusic* Library

4.1 *OpenMusic*, a Graphical Music Processing Language

OpenMusic was originally designed, by Gérard Assayag (1999) and Carlos Agon, for Computer Assisted Composition. But since it was aimed at offering the most general tools as possible for music processing, it may be used for any processing of abstract musical representations, and even for non musical ones. This software offers a highly object-oriented visual environment for Common LISP programming completed by numerous abstract music processing tools. Objects are symbolized by icons that may be dragged and dropped all around. Most operations are then performed by dragging an icon from a particular place and dropping it in another. A lot of classes implementing musical data and behavior are provided. They are associated with graphical editors and may be visually subclassed by the user to meet specific needs. Different representations of a musical process are handled, which include common notation, MIDI piano-roll, sound waveforms, break-point functions. A *patch* is a place where objects -

⁴ For more details about the implementation, and description of additional mechanisms, see Lartillot (2003).

repetitions, as we have seen, generally feature local temporal and intervallic distortions. Even recorded performance can be analyzed in this exact way. In fact, the motivic dimension may ignore enharmonic information. In the *Prelude*, for instance, the interval between the two bass notes F# (bar 22) and Ab (bar 23), “motivically speaking”, simply sounds as a 2-half-tone-upward step, be it a diminished third or an exact second.

In current version, results are displayed as a list of texts that is not easy to understand. That is why this library is also provided with some basic tools for selecting and displaying longest patterns, most frequent patterns, or most pertinent patterns, where pertinence is a product of length and frequency.

When asking the pattern classes achieving the highest degree of pertinence in the beginning of Bach’s *Prelude* in C, we obtain the 8-note pattern [Figure 20, first line] and 3-note patterns that are repeated inside the 8-note pattern itself [second line]. Moreover, the beginning of the two last bars are considered as prefixes of this 8-note pattern [Figure 20b, first line], and in the same time as two new similar patterns [second line]. Such interesting results, discovered by a system that has no access to any prior cultural knowledge, has never been obtained by any previous algorithm. The mechanisms of meta-pattern discovery presented in paragraph 2.6 are being implemented now and results will be soon available.



Figure 20a. Most pertinent patterns of the beginning of *Prelude* in C Major, BWV 846, by Bach.




Figure 20b. Patterns at the ending of the *Prelude*.

When asking the most pertinent pattern classes in the beginning of the *Fifth Symphony* by Beethoven, we obtain the 4-note pattern [Figure 21a-b, second line] and the encapsulation of 3 times the 4-note patterns [Figure 21b, third line]. Note however that a non-pertinent occurrence has been found (dotted lined). Non-pertinent encapsulation of 2 times the 4-note patterns [first line] has also been considered.




Figure 21a. Most pertinent patterns of the beginning of the *Fifth Symphony* by Beethoven.




Figure 21b. continued.

When asking the most pertinent pattern classes in the beginning of the *Rondo alla turca* by Mozart, we obtain the 5-note pattern [Figure 22a, second line]. Its prefix is also displayed [third line] because little patterns of four descending notes [Figure 22b] have been considered as occurrences of this pattern class. The algorithm is not well prepared to face polyphonic pattern. That is why we observe some failures in pattern detection [Figure 22a, first and fourth lines]. Finally, the long pattern [fifth line] pertinently encapsulates the repetitions of the 5-note pattern, but cannot then integrate the following chords, due to the polyphony limitations.




Figure 22a. Most pertinent patterns of the beginning of the *Rondo alla turca* by Mozart.




Figure 22b. continued (only the first three lines).

Such results may appear very simple and evident, compared to the complexity of the conceptual framework that has been developed for this purpose. But it has to be remarked that such results are very interesting for

Artificial Intelligence researches, and also for the understanding of the cognitive mechanisms of music perception. Moreover, now that the algorithm performance starts corroborating with expecting behavior, a little refinement of the model should offer more subtle results [Figure 11]. A corpus of several pieces may be analyzed altogether, in order to find common patterns.

4.3 Limitations and Future Works

The algorithm may induce non-pertinent knowledge too, and a significant amount of analytical concepts are still outside of the score of the model. In fact, it has to be acknowledged that this research is still in a very early phase, and that the library remains as a rough prototype with innumerable limitations. Many investigations need to be undertaken before obtaining a valuable tool for computer aided musicology.

There may be, inside patterns, “enclaves” of foreign notes not really belonging to these patterns. Patterns may also features transitory states, such as passing notes or appoggiatura. More generally, patterns may be included in a polyphonic flow. If all this flow is represented as a single totally ordered sequence, patterns representations, here also, feature enclaves. Such problem has already been tackled (see David Meredith et al. 2002), but uniquely for exact repetition. Chords should also be taken into consideration and patterns of chords should be discovered.

Then an interface has to be designed, enabling a browsing inside the score and the discovered structures. In a long term, such approach may go beyond pattern and catch higher-level concepts. A project of automatic music theory discovery may also be envisaged.

5 Conclusion

The new approach of musical pattern discovery we proposed, based on a modeling of cognitive mechanisms of music perception, leads to promising results. Most of the discovered structures correspond to basic patterns effectively perceived by human listeners. With some improvements, this algorithm should also be able to detect more subtle patterns that are less easily discriminated by human listeners, but that participate to the complex flow of implicit reasoning, which constantly submerges the mind of an experienced listener when enjoying musical experience.

Acknowledgments

My doctorate project is supervised by Emmanuel Saint-James (LIP6, Paris 6) and Gérard Assayag (Musical Representations Team, Ircam). Lots of ideas stem from stimulating discussions with my colleague Benoit Meudic. Up-to-date information about this project is available at following webpage address: www.ircam.fr/equipes/repmus/lartillot

References

- Andreatta, Moreno, and C. Agon, "Algebraic Methods in Music: Theory and Implementation," *17th International Congress of the International Musicological Society* (2002).
- Assayag, Gérard, C. Rueda, M. Laurson, C. Agon, and O. Delerue, "Computer Assisted Composition at Ircam : PatchWork and OpenMusic," *Computer Music Journal*, 23/3 (1999), 59-72.
- Cambouropoulos, Emilios, "Towards a General Computational Theory of Musical Structure." Ph.D. Thesis, University of Edinburgh, 1998.
- Conklin, Darrell, and I. Witten, "Multiple Viewpoint Systems for Music Prediction," *International Journal of New Music Research*, 24/1 (1995), 51-73.
- Cook, Nicholas. *A Guide to Musical Analysis*. London: Dent, 1987.
- Crawford, Tim, C.S. Iliopoulos, and R. Raman, "String-Matching Techniques for Musical Similarity and Melodic Recognition," *Computing in Musicology* 11 (1997-1998), 73-100.
- Dannenber, Roger, "Listening to "Naima": An Automated Structural Analysis of Music from Recorded Audio," *Proceedings of the International Computer Music Conference* (2002), 28-34.
- Dowling, W.J., "Scale and contour: Two components of a theory of memory for melodies," *Psychological Review* 85/4 (1978), 341-354.
- Lartillot, Olivier, "Musical Analysis by Computer Following Cognitive Model of Induction of Analogies," *Proceedings of the International Computer Music Conference* (2002), 20-27.
- Lartillot, Olivier, "Discovering Pertinent Pattern Through Perceptual Heuristics: Why and How," *Proceedings of the International Conference on Music Information Retrieval* (2003), to appear.
- Lerdahl, Fred, and Ray Jackendoff. *A Generative Theory of Tonal Music*. Cambridge, MA: MIT Press, 1983.
- Lindsay, Adam, "Using Contour as a Mid-Level Representation of Melody", Masters Thesis, MIT, 1994.
- Meredith, David, K. Lemström, and G. A. Wiggins, "Algorithms for Discovering Repeated Patterns in Multidimensional Representations of Polyphonic Music," *Journal of New Music Research* 31/4 (2002), 321-345.
- Meyer, Leonard B., *Emotion and Meaning in Music*. Chicago: The University of Chicago Press, 1956.
- Narmour, Eugene, *The Analysis and Cognition of Basic Melodic Structures: The Implication-Realization Model*, The University of Chicago Press, 1990.
- Nattiez, Jean-Jacques, *Music and Discourse: Towards a Semiology of Music*. Princeton: Princeton University Press, 1990.

- Reti, Rudolph R., *The Thematic Process in Music*. New York: Macmillan, 1951.
- Ruwet, Nicolas, "Methods of Analysis in Musicology," *Music Analysis* 6/1-2 (1987), 4-39.
- Selfridge-Field, Eleanor, "Conceptual and Representational Issues in Melodic Comparison," *Computing in Musicology* 11 (1997-1998), 3-64.
- Temperley, David, *The Cognition of Basic Musical Structures*, Cambridge, MA: MIT Press, 2001.