



HAL
open science

Discovering Musical Patterns through Perceptive Heuristics.

Olivier Lartillot

► **To cite this version:**

Olivier Lartillot. Discovering Musical Patterns through Perceptive Heuristics.. ISMIR 2003, Oct 2003, Baltimore, United States. pp.1-1. hal-01161139

HAL Id: hal-01161139

<https://hal.science/hal-01161139>

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.

Discovering Musical Patterns through Perceptive Heuristics

Olivier Lartillot

IRCAM -Centre Pompidou

4, Place Igor-Stravinsky,

75004 Paris, France

Olivier.Lartillot@ircam.fr

Abstract

This paper defends the view that the intricate difficulties challenging the emerging domain of Musical Pattern Discovery, which is dedicated to the automation of motivic analysis, will be overcome only through a thorough taking into account of the specificity of music as a perceptive object. Actual musical patterns, although constantly transformed, are nevertheless perceived by the listener as musical identities. Such dynamical properties of human perception, not reducible to geometrical models, will only be explained with the notions of contexts and expectations. This paper sketches the general principles of a new approach that attempts to build such a general perceptual system. On a sub-cognitive level, patterns are discovered through the detection, by an associative memory, of local similarities. On a cognitive level, patterns are managed by a general logical framework that avoids irrelevant inferences and combinatorial explosion. In this way, actual musical patterns that convey musical significance are discovered. This approach, offering promising results, is a first step toward a complete system of automated music analysis and an explicit modeling of basic mechanisms for music understanding.

1 Introduction

Musical Pattern Discovery (MPD) is an emerging discipline, which aims at offering automated motivic analyses of musical scores. Nowadays, however, no computer system is able to complete in a relevant way the demanding task of MPD. In this paper, we suggest a new approach that may offer an answer to fundamental problems arisen in this discipline, and that could therefore solve a certain number of difficulties.

We introduce a general methodology whose intuition stems from a mimicking of human perception. We will defend such a position through a critical overview of current approaches in

MPD. We suggest a characterization of musical pattern that takes into account the fundamental notions of repetition, sequence and similarity in the musical context of temporal perception. Following Hofstadter's propositions (1995), we propose to discriminate between sub-cognitive and cognitive mechanisms. We sketch the general principles of a computational model that attempts to pursue such methodology. The model has been implemented and offers first promising results.

2 Pattern Discovery in Music Analysis

2.1 The Dimensions of Music Analysis

Music may be considered as a particularly complex "language" that conveys information or meaning along different dimensions.

Melody is one of its most salient aspects, since it is traditionally ruled by mnemonic constraints that help recognition and learning. Accompaniment too features salient "syntactic" or *motivic* aspects. *Motivic analysis* may therefore focus on all these aspects.

Rhythm is another essential aspect of music, that may be considered separately, but that also play an important part in melodies and motives.

Metric is a global vision of temporal structures that is usually dependent on stylistic knowledge.

Moreover, music is fundamentally ruled by the "grammatical" rules of harmony. Such rules participate more or less in the comprehension of music by average listeners, but only in a rather implicit way. This higher-level symbolic discourse, which can be objectified by educated listeners, also offers salient aspects, such as cadenza. Its analysis may therefore share some similar aspects with motivic analysis, such as the inductive methods of patterns discovery and recognition.

Finally, the construction of the musical work ends up in a global formal structure, that can be considered in a very schematic way — such as *A-B-A* — or more carefully. But as the formal structure is a simple consequence of the material construction, its detailed description may be simply revealed by a thorough motivic analysis.

Our approach will focus on the design of an automated system for motivic analysis, without any reference to harmony or style.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. © 2003 The Johns Hopkins University.

2.2 Music Groupings

The significant structures discovered by music analysis naturally consist of sets of notes. Hence analyzing means grouping local notes in the score, and relating these groups one to others and to external concepts. These groupings are generally guided by three principal criteria:

2.2.1 Style-Based Groupings

Metrical and harmonic descriptions of music may induce a specific way of segmenting the musical discourse, especially with the help of stylistic norms, which may for instance constrain pattern length (Lerdahl and Jackendoff, 1983). As our approach will not rely on harmonic and stylistic knowledge, such heuristics will not be considered.

2.2.2 Local Boundaries

The musical surface may be segmented by *local* discontinuities, in particular between notes that are particularly distant in time or in pitch dimensions (Lerdahl and Jackendoff, 1983), or when the melodic contour changes sign, i.e. when melody first grows and then decreases, for instance. Discontinuities along other parameters¹ such as dynamics, accents or instrumentation may contribute to these local segmentations (Cambouropoulos, 1998).

2.2.3 Repetition

Finally, a musical motive is traditionally defined as a set of notes that is *repeated* several times throughout the score (Nattiez, 1990). Such criterion conflicts with the previous one. Indeed, a motive may feature contrastive steps. In Figure 1, the first leap between the first two notes, although segmenting the local discourse, is the important element that characterizes the beginning of the motive itself. Instead of unifying what is constant and alike, a motive rather unifies what is contrastive and different, or more precisely: what is similarly contrastive in different places of the work.



Figure 1: A motive (solid lines) may feature contrastive steps and therefore contradicts local segmentations (dashed lines).

Our approach will uniquely focus on repetition, since the concept of repeated motives seems more developed in traditional music theories than the notion of local boundaries. Nevertheless, the process of repetition discovery itself is ruled by local constraints too. Indeed, repeated patterns are detected only when their successive notes are sufficiently close, and when their borders contrasts sufficiently with the outer environment, in particular because of temporal or pitch discontinuities. For instance, a pattern that is strictly included inside an ascending movement (first dashed-lined grouping in Figure 2) cannot be related to its exact repetition (second dashed-lined grouping) since it has not been pre-segmented by local boundaries first.



Figure 2: Exactly or approximately repeated pre-segmented patterns (solid lines) are perceived, whereas exactly repeated but non-pre-segmented ones (dashed lines) are not perceived.

2.3 Towards Conceptual Pattern Discovery

The previous illustrated discussion implies that repetition of successions — that may be discovered through string algorithms (Smith and Medina, 2001) —, although the core characteristic of musical motives, is not a sufficient condition for their exhaustive determination. Other aspects contribute to this idea. For instance, a series of successive repetitions of a single elementary pattern features numerous geometrical repetitions that are not relevant.




Figure 3: Irrelevant geometrical patterns induced by the successive repetition of an elementary pattern.

Cambouropoulos (1998) proposes to solve such paradox by selecting classes of patterns where overlapping is minimal. This means that all possible patterns are computed first, and that certain classes of patterns are then selected according to global factors, with no respect to each particular case. Such a statistical approach, which may offer satisfying results for simple examples, cannot be applied to actual score where classes of patterns may be relevant at some places and not at others.

This difficulty may be solved by considering the fact that patterns are conceptually inferred during the incremental listening of the piece. For each new note, the set of inferences currently in process constitutes a context, which induces constraints upon the candidates for new inferences. In this way, irrelevant inferences described in previous examples may be avoided.

2.4 Approximate Repetitions

Musical motives are commonly subject to numerous kinds of transformations, such as transposition (pitch translation), time stretching, distortions due to scale constraints, and thematic developments.

2.4.1 Related Works

One popular approach for detecting patterns repetition is founded on self-similarity matrix (Dannenberg, 2002). Music is decomposed into local segments $s = 1, \dots, N$; similarity

¹ In a first approach, these parameters will not be considered.

distances are measured between all possible pairs (i, j) of local segments and stored in a matrix at position (i, j) . Then patterns may be discovered along lines of the matrix that are (approximately) parallel to the first diagonal. Such approach, by implicitly assuming that one necessary condition for pattern similarity is the similarity of their elementary constituents at each respective position, cannot identify transformed patterns, even simply transposed ones. Actually, what characterizes a pattern is less its intrinsic composition than the formal properties that are shared by its different repetitions. One particular property is the local relationships between notes, which is not actually taken into account in methods based on self-similarity matrix.

In order to handle such music plasticity, multiple parallel “viewpoints” (Conklin and Anagnostopoulou, 2001) or “level of abstraction” (Dannenberg, 2002) of the pattern are considered. For instance, transposed patterns may be identified by expressing pitch parameters relatively to a reference point associated to each pattern, such as their respective first note. A simpler heuristic consists of preferring pitch intervals between successive notes to absolute pitches. However, such methods cannot handle local distortion inside patterns. In this case, more relative viewpoints may then be considered. One commonly used description is the contour, which describes the direction of interval between successive notes (downward, upward, or constant). The trouble is, contour-based repeated pattern discovery algorithms easily produce irrelevant results, since lots of false positive patterns may be discovered.

2.4.2 Memories and Representations

Dowling and Harwood (1986) have suggested that contour is the minimal degree of similarity between patterns, but also that contour description is relevant for short-term memory (STM) retrieval whereas pitch interval description is necessary for long-term memory (LTM). Indeed, the very small scale of STM and its very active presence in our mind makes it easy to relate its content to present perception in very loose terms, such as contour representation. On the contrary, due to its large size, LTM should be searched through precise queries, such as pitch interval representations.

That is why it may be interesting to divide the pattern discovery process into two subtasks, one for STM and one for LTM. This paper will only focus on the LTM part. We suggest in paragraph 3 an improvement of Dowling’s suggestion. It may be interesting to limit his constraint on the query exactness to only a *prefix* of it. Indeed, if a prefix of the query is sufficient for the retrieval of one particular context in LTM, the remaining of the retrieved pattern may simply be compared to the remaining of the query along contour representation.

Other factors may influence pattern discovery. For instance, patterns that are highly expected, for instance after having just been successively repeated several times before, as will be shown in paragraph 4.6, can be detected even when highly distorted. Moreover, patterns featuring a particular structure may be retrieved each time such structure is discovered again.

All this supports the claim that pattern discovery needs perceptual considerations, related in particular to time and memory (Lartillot and Saint-James, 2003).

3 Subcognitive Mechanisms of Pattern Discovery

3.1 Local contexts

3.1.1 Associative Memory

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 explicitly defined before.

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.

3.1.2 Interval Distances

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'_1 and n'_2 be four notes whose respective pitches are p_1, p_2, p'_1 and p'_2 and respective time onset o_1, o_2, o'_1 and o'_2 . We propose, in a first approach, to formalize the perceptual distance between the two intervals (n_1, n_2) and (n'_1, n'_2) as a weighted product of a pitch interval distance and an inter-onset distance:

$$D_2((n_1, n_2), (n'_1, n'_2)) = (\text{abs} [(p_2 - p_1) - (p'_2 - p'_1)] + 1) * (\max [(o_2 - o_1) / (o'_2 - o'_1), (o'_2 - o'_1) / (o_2 - o_1)])^{0.7}$$

We may remark that in music, pitch intervals are subtracted, whereas inter-onsets are divided.

3.2 Initiation Phase: Discovering similar contexts

Each 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 pitch interval value the set of its associated occurrences in the part of the score that has already been analyzed. For any current local interval (n_1, n_2) of pitch intervals $(p_2 - p_1)$, any *similar* old local interval (n'_1, n'_2) belongs to those stored in the hash-table at a hash value $(p'_2 - p'_1)$ equal or *similar* to $(p_2 - p_1)$, that is:

$$\text{abs} [(p'_2 - p'_1) - (p_2 - p_1)] < \square,$$

Similar hash-tables could be added for other interval parameters such as inter-onset values, in order to retrieve in particular similar rhythmic patterns.

For each similar old interval (n'_1, n'_2) , its previous interval (n'_0, n'_1) is now considered, and compared to the previous interval (n_0, n_1) of (n_1, n_2) . If (n'_0, n'_1) and (n_0, n_1) are also similar, then (n_0, n_1, n_2) and (n'_0, n'_1, n'_2) are considered as similar patterns.




Figure 4: Similarity between intervals is found through the hash-table of local intervals (up). When two intervals (n_1, n_2) and (n'_1, n'_2) are similar and preceding intervals (n_0, n_1) and (n'_0, n'_1) are similar too, then a pattern is inferred (down).

Moreover, when (n_1, n_2) and (n'_1, n'_2) are first compared, if additional facts contribute to their equivalence — for instance $p_1 = p'_1$ —, then the inference of the repetition of patterns (n_1, n_2) and (n'_1, n'_2) may be immediately inferred without considering their previous patterns. If on the contrary similarity between (n_0, n_1, n_2) and (n'_0, n'_1, n'_2) is not high enough, maybe previous intervals can be considered once again in a recursive way.

3.3 Extension Phase: Prolonging contexts

Once the similarity between the beginnings of two patterns have been detected, the inference of the similarity between the continuations of each pattern is far easier to compute, since there is no need to discover new contexts any more. Therefore, continuations can be compared only along their contour parameters.




Figure 5: Once a pattern is inferred, respective following intervals can be easily compared along their contour.

3.4 Local Pre-Segmentation

As explained in paragraph 2.2.3, repeated pattern can be discovered only when they sufficiently contrast to their surrounding at their borders. Their should be a *discontinuity* at the beginning of current candidate pattern (n_0, n_1, \dots) and at the beginning (or at the ending) of past candidate pattern $(n'_0,$

$n'_1, \dots)$. We propose to sketch roughly such idea as follows: there is a *discontinuity* between the note n_- preceding the pattern (n_0, n_1, \dots) and the pattern itself when either:

- the contour changes between intervals (n_-, n_0) and (n_0, n_1) , from up to down, from constant to up, etc.
- there is a significant temporal gap before the pattern: $(o_0 - o_-) / (o_1 - o_0) > \square_b$
- there is a significant pitch interval before the pattern: $(p_0 - p_-) / (p_1 - p_0) > \square_b$

In current implementation, \square_b and \square_b are fixed to 2, but the value should perhaps be lower and dependent on context.

4 Logical Rules of Pattern Processing

4.1 Incremental Conceptual Building on the Score

The discovery of single patterns, fundamental mechanisms of which have been proposed in previous section, though one essential characteristic of musical pattern perception, does not suffice to take into account human capabilities. Logical principles have to be added in order to make sure that the cognitive system infer relevant knowledge and do not face combinatorial explosion.

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.

The score is progressively scanned from the beginning to the end, each new note triggers the calling of a main routine that carries out a determinate set of analysis operations.

4.2 Pattern Classes and Pattern Occurrences

When a pattern is repeated several times, each occurrence should not be related to each other, since this would lead to an undesired complex relationship network. What has to be taken into account here is the concept of *pattern classes* (PC) and *pattern occurrences* (PO), also called “clustering” (Dannenberg, 2002). Multiple repetitions of a pattern should therefore simply be considered as occurrences of a single PC.

A pattern is an approximately repeated succession of notes. Therefore a PO is a set of notes in the score, whose succession is similar to the succession defined in the PC characteristics. Such property may be formalized through graph theory.

We propose to model a pattern as a chain of states, where successive states correspond to successive notes of the pattern, as in Figure 6. In this way, a PC is a chain of states, called *pattern class chain* (PCC), where each state represents the shared characteristic of the associated note, and where each transition between two successive states represents an interval between two successive notes.

Each PO is also a chain of states, named *pattern occurrence chain* (POC), where each state interfaces a particular note in the score with its corresponding state in the associated PC.




Figure 6: The two POCs (black circles) interface notes in the score with the corresponding states in the PCC (white circles).

4.3 Discovering Further Occurrences

Once a PC has been discovered, its further POs should all be subsumed under the same PC. The distinction between context discovery and context prolonging still prevails.

During *initiation phase*, when a similarity has been discovered between two different contexts, and before deciding to create any new PC, we have to make sure that such context does not already exist in the beginning of one of the set of all discovered PCs. If there does exist such pattern class, a new PO simply associates the new discovered context with the retrieved PC.

The *extension phase* of PO discovery significantly differs from PC discovery. Since the beginning of currently discovered PO is already associated to a PC, each of its successive candidate continuations may simply be compared to the successive continuations along the PC. In this case, current PO does not need to be compared to old occurrences.

4.4 Avoiding Redundancy

We have to constantly check that each new pattern occurrence discovery was not already induced through a previous mechanism that led to the same result. The incremental and logical thinking that builds human perception of music is ruled by fundamental principles, which are necessary for a coherent process. For example, every time a sequence is considered as an occurrence of a pattern, every suffix could itself be considered as occurrence of other pattern class, for simple mathematical reasons. But cognitively speaking, such inferences are irrelevant, since they do not correspond to inference human makes when listening to music. This is due to the fact that the first longest pattern was sufficient to explain the phenomenon, and that further inferences of suffixes would only infringe a clear analysis of the score. That is why suffix of pattern or subsequences of sequences should not be explicitly represented. (Dannenberg, 2002; Pienimäki, 2002). Our formalization through graph theory enables to easily implement such rules.




Figure 7: Suffixes (bottom POCs and PCC) cannot be considered as new PCs, since the associated notes of each of

its POCs already belongs to the POCs of the whole pattern (top POCs and PCC).

4.5 Parallel Memory vs. Sequential Computer

The whole configuration of relationships between notes, PCs, and POs, forms a particularly complex and interdependent system. Such vision is largely influenced by Holland et al.'s (1989) cognitive model of induction. All the operations undertaken during the analysis are in fact triggered or inhibited by the configuration of the network and the prior rules. The trouble is, the implementation of such a parallel system has to cope with the sequential architecture of computers. For this purpose, the operations has to be very precisely ordered, else numerous artifacts may produce irrelevant results and combinatorial explosions.

For each new note considered by the main routine, each possible PO concluded by previous note is a candidate for three successive operations:

4.5.1 Pattern Occurrence Extension

Current new note may be associated to one possible continuation of the PO, such that this continuation is already described in the PC. Candidates are considered in a decreasing order of similarity. Minor candidates that are considered as negligible compared to the most salient ones are discarded. In this way, current note may induce poorly perceptible patterns only when it cannot be related to particularly salient patterns.

4.5.2 Pattern Class Extension

If previous condition *does not occur*, the eventual extension of the PO with current note into a new PC is checked. This extension should not be already inferred by current configuration. As previously, negligible candidates are discarded.

4.5.3 Pattern Class Initiation

Finally, in any case, new pattern initiations are attempted. As for pattern extension, these new patterns should not be already deduced by the previously inferred POs. In particular, as shown in paragraph 4.4, simple suffixes of patterns cannot be considered as new patterns.

If the past candidate (n'_0, n'_1, n'_2) already initiate another PC, then current candidate (n_0, n_1, n_2) is simply associated to a new PO of this PC.

Finally, for all three trials, it is also necessary to make sure that every PC or PO in process of creation does not already exist. The inference of duplicate concepts may indeed provoke disastrous effects such as combinatorial explosions.

4.6 Pattern Association

Music features multiple levels of pattern descriptions. Notes may indeed belong to several possible patterns in parallel. When discovering another PO of a PC that was associated to another PC, we may expect to retrieve the same association between these PCs. Therefore, pattern association may induce pattern *expectation*.

For this purpose, the association between patterns may be directly represented on the PCs themselves. Every time a note of a PO is associated to another PO, on the corresponding state

in each PC is associated a new pattern associated to the other PC. Such pattern association discovery induces a pattern association expectation rule, stating that every time a new PO of such PC is discovered, possible associated PCs are also expected. Such expectation may be formalized through the initiation of new hypothetical PO where only the first state is represented, waiting for further extension.

For instance, a pattern may include another sub-pattern. In Figure 8, the 8-note PO features a repetition of two 3-note POs. Therefore, on the 8-note PC itself is added two 3-note POs.




Figure 8: The 3-note POCs (up) that are included inside the 8-note POC are represented directly on the 8-note PCC (down).

In this way, during the discovery of a new 8-note POC, corresponding 3-note POCs are also expected (right).

4.7 Successive Repetitions

Music usually features successive repetitions of a same PC. As explained in paragraph 2.3, if no new mechanisms were added, the system would consider each possible concatenation of these successive patterns as a new pattern. These inferences, not corresponding to human judgments and leading to combinatory explosion, should be forbidden.

It should 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 PC. This means that, in the extended PC, the last state is linked to the first state. The idea of pattern cycling is therefore explicitly represented.

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




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

4.8 Meta-Pattern of Patterns

The succession of POs may be examined exactly the same way that the previously considered succession of simple notes, provided that these POs belong to the same PC. For this purpose, a concept of interval between two successive POs, that generalizes the traditional notion of interval between two successive notes, is defined. The parameters of an interval between notes were pitch intervals and inter-onset. In a similar way, we define the parameters of an interval between POs as:

- the pitch intervals and inter-onsets between the respective *last notes* of each PO. In this way, the interval between notes of previous positions in each PO can be obtained by considering the interval between POs prefixes.
- the difference between pitch intervals between *the respective last two notes* of each PO. Idem for inter-onsets. Such parameter, that may be considered as a meta-interval between intervals, represent the evolution of the interval value between the two POs.




Figure 10: In a generalized interval between pattern occurrences, notes of same rank are compared.

As intervals between notes were stored in a hash-table, intervals between POs are stored in a new hash-table dedicated to the associated PC. Each time a new PO is inferred (be it the initiation of a new pattern or the extension of a previous one), the meta-interval between the previous PO of same PC and current PO is stored in the PC hash-table.

In this way, the meta-pattern discovery process is a straightforward generalization of pattern discovery. Instead of writing new functions dedicated to meta-pattern discovery, we propose to directly consider simple note as a degenerate case of PO, and the concept of note as a special PC. Hence the main routine that is called for each new note is also recursively called by sub-routines of the routine itself, each time a new PO is inferred.

When two intervals between POs of same PC are similar, a meta-PC is initiated, with the two corresponding meta-POs.




Figure 11: Initiation of a new meta-PC.

The main difficulty of the meta-generalization is that the elements of a meta-pattern may themselves be extended. If current PO is extended, and if previous PO of current meta-PO and previous meta-PO(s) can also be extended, then these

meta-POs will contain these extended POs instead of the initial POs.




Figure 12: Extension of the element of the meta-PC.

Meta-POs may extend their POs even when intervals between extended POs are no more similar. For instance in Figure 12, although the second note of successive POs (E-F-F-E) do not form any pattern (or maybe a mirrored one), the initial 1-note POs are extended to 2-note POs, since the overall succession of these POs are actually perceived as meta-patterns.

On the contrary, intervals between extended POs may sometimes contribute to the initial meta-pattern. For instance, the third note of POs (G-G-G-G) also form a same pattern (G-G, G-G).

Exact repetition of POs may be considered as generalized “unison”. Moreover, this meta generalization may be operated recursively and meta-pattern of meta-pattern may be considered too!

5 Current Results

5.1 Implementation

These algorithms have been implemented in Common Lisp language, as a library of *OpenMusic* (Assayag et al., 1999), a graphical programming language dedicated to the computation of symbolic representations of music. In current version of this library, called *OMkanthus*, these results are displayed as a list of texts that is not easy to understand. That is why this library is provided with some basic tools for selecting and displaying longest patterns, most frequent patterns, or most significant patterns, where pertinence is a product of length and frequency.

5.2 Results

The analysis of Bach’s *Prelude in C*, BWV 846, generates all the occurrence of the 8-note pattern. The whole *Prelude* (except the last two bars) is represented as a perpetual repetition of a meta-pattern of exact repetition of two 8-note pattern occurrences. Meta-pattern of meta-pattern are also discovered, such as transpositions of sequences of two and four bars at several different places.




Figure 13: Each bar of the Bach’s *Prelude* is a meta-pattern of an exact repetition of two 8-note patterns.




Figure 14: Two transposed occurrences of a meta-patterns of two meta-patterns of 8-note patterns.




Figure 15: Two transposed occurrences of a meta-patterns of four meta-patterns of 8-note patterns.

The last bars of the *Prelude* feature the first two notes of the constantly expected 8-note patterns, but with a new continuation that constitutes a new pattern class.



Figure 16: The last bars feature two occurrences of a new pattern, which is developed from the first two notes of 8-note patterns.

Irrelevant patterns may sometimes be found too. This model should be considered more as a very experimental prototype that attempts to simulate some aspects of pattern perception, than as a complete and robust automated music analyzer.

6 Future Works

6.1 Integrating Contour Representation

As explained previously, the contour representation is highly relevant for short-term memory retrieval. A specific pattern discovery module dedicated to contour representation would enable to easily represent local relationships between similar gestures.

6.2 Large-Scale Intervals

Each successive note of the currently heard repetition is tentatively related with a possible successive note of the previously heard occurrence. 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 Figure 17, 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 17: 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.

6.3 Towards Polyphony

There may be, inside patterns, “enclaves” of foreign notes not really belonging to these patterns. Patterns may also feature transitory states, such as passing notes or appoggiatura.

More generally, patterns may be included inside a polyphonic flow. If all this flow is represented like a single totally ordered sequence, patterns representations, here also, feature enclaves. Such problem has already been tackled by Meredith, Lemström and Wiggins (2002) but for exact repetition.

Chords should also be taken into consideration and patterns of chords should be discovered.

7 Conclusion

In this paper, we have proposed a new approach of musical pattern discovery based on a modeling of cognitive mechanisms of music perception. This strategy leads to promising results. 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.

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.

Acknowledgements

My doctorate project is supervised by Emmanuel Saint-James (LIP6, Paris 6) and Gérard Assayag (Musical Representation Team, Ircam). Lots of ideas arise from very stimulating discussion 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

- Assayag, G., Rueda, C., Laurson, M., Agon C. & Delerue, O. (1999). Computer Assisted Composition at Ircam : PatchWork and OpenMusic. *Computer Music Journal*, 23(3), 59-72.
- Cambouropoulos, E. (1998). *A Generative Theory of Tonal Music*. Ph.D. dissertation, Edinburgh Univ.
- Conklin, D. & Anagnostopoulou, C. (2001). Representation and Discovery of Multiple Viewpoint Patterns. In *Proceedings of the International Computer Music Conference*, (pp.479-485). San-Francisco: ICMA.
- Dannenberg, R. (2002). Listening to “Naima”: An Automated Structural Analysis of Music from Recorded Audio. In *Proceedings of the International Computer Music Conference*, (pp.28-34). San-Francisco: ICMA.
- Dowling, W. J. & Harwood, D. L. (1986). *Music Cognition*. Orlando, FL: Academic Press.
- Hofstadter, D. (1995). *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*. New-York: Basic Books.
- Holland, J. H., Holyoak, K. J., Nisbett, R. E. & Thagard, P. R. (1989). *Induction: Processes of Inference, Learning, and Discovery*, Cambridge, MA: MIT Press.
- Lartillot, O. & Saint-James, E. (2003). Automating Motivic Analysis through the Application of Perceptual Rules. *Computing in Musicology*, 13, to appear.
- Lerdahl, F. & Jackendoff, R. (1983). *A Generative Theory of Tonal Music*. Cambridge, MA: MIT Press.
- Meredith, D., Lemstrom, K. & Wiggins, G. A. (2002). Algorithms for discovering repeated patterns in multidimensional representations of polyphonic music. *Journal of New Music Research*, (31)4, 321-345.
- Nattiez, J.-J. (1990). *Music and Discourse: Towards a Semiology of Music*. Princeton, NJ: Princeton University.
- Pienimäki, A. (2002). Indexing Music Databases Using Automatic Extraction of Frequent Phrases. In *Proceedings of the International Conference on Music Information Retrieval*, (pp.25-30). Paris: Ircam.
- Smith, L. & Medina, R. (2001). Discovering Themes by Exact Pattern Matching. In *Proceedings of the International Conference on Music Information Retrieval*, Bloomington, IN.