



Aalborg Universitet

AALBORG UNIVERSITY
DENMARK

In-depth motivic analysis based on multiparametric closed pattern and cyclic sequence mining

Lartillot, Olivier

Published in:
15th International Society for Music Information Retrieval Conference

Publication date:
2014

Document Version
Accepted author manuscript, peer reviewed version

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Lartillot, O. (2014). In-depth motivic analysis based on multiparametric closed pattern and cyclic sequence mining. In H-M. Wang, Y-H. Yang, & J. H. Lee (Eds.), 15th International Society for Music Information Retrieval Conference

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- ? Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- ? You may not further distribute the material or use it for any profit-making activity or commercial gain
- ? You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

IN-DEPTH MOTIVIC ANALYSIS BASED ON MULTIPARAMETRIC CLOSED PATTERN AND CYCLIC SEQUENCE MINING

Olivier Lartillot

Aalborg University, Department of Architecture, Design and Media Technology, Denmark
olartillot@gmail.com

ABSTRACT

The paper describes a computational system for exhaustive but compact description of repeated motivic patterns in symbolic representations of music. The approach follows a method based on closed heterogeneous pattern mining in multiparametrical space with control of pattern cyclicity. This paper presents a much simpler description and justification of this general strategy, as well as significant simplifications of the model, in particular concerning the management of pattern cyclicity. A new method for automated bundling of patterns belonging to same motivic or thematic classes is also presented.

The good performance of the method is shown through the analysis of a piece from the JKUPDD database. Ground-truth motives are detected, while additional relevant information completes the ground-truth musicological analysis.

The system, implemented in Matlab, is made publicly available as part of *MiningSuite*, a new open-source framework for audio and music analysis.

1. INTRODUCTION

The detection of repetitions of sequential representations in symbolic music is a problem of high importance in music analysis. It enables the detection of repeated motifs and themes¹, and of structural repetition of musical passages.

1.1 Limitation of previous approaches

Finding these patterns without knowing in advance their actual description is a difficult problem. Previous approaches have shown the difficulty of the problem related to the combinatorial explosion of possible candidate patterns [2]. Some approaches tackle this issue by generating a large set of candidate patterns and applying simple global heuristics, such as finding longest or most frequent patterns [3,8]. Similarly, other approaches base the search for patterns on

¹ Here motif and theme are considered as different musicological interpretations of a same pattern configuration: motifs are usually shorter than themes.

general statistical characteristics [5]. The problem is that there is no guarantee that this global filtering leads to a selection of patterns corresponding to those selected by musicologists and perceived by listeners.

1.2 Exhaustive mining of closed and cyclic patterns

In our research, we endeavour to reveal the factors underlying this structural explosion of possible patterns and to formalise heuristics describing how listeners are able to consensually perceive clear pattern structures out of this apparent maze. We found that pattern redundancy is based on two core issues [6]:

- *closed pattern* mining: When a pattern is repeated, all underlying pattern representations it encompasses are repeated as well. In simple string representation, studied in section 2², these more *general* patterns correspond to prefixes, suffixes and prefixes of suffixes. The proliferation of general patterns, as shown in Figure 1, leads to combinatorial explosion. Restricting the search to the most specific (or “maximal”) patterns is excessively selective as it filters out potentially interesting patterns (such as CDE in Figure 1), and would solely focus on large sequence repetitions. By restricting the search to *closed* patterns – i.e., patterns that have more occurrences than their more specific patterns –, all pattern redundancy is filtered out without loss of information. [6] introduces a method for exhaustive closed pattern mining.
- *pattern cyclicity*: When repetitions of a pattern are immediately successive, another combinatorial set of possible sequential repetitions can be logically inferred [2], as shown in Figure 2. This redundancy can be avoided by explicitly modelling the cyclic loop in the pattern representation, and by generalising the notion of closed pattern accordingly.

By carefully controlling these factors of combinatorial redundancy without damaging the non-redundant pattern information, the proposed approach in [6] enables to output an exhaustive description of pattern repetitions. Previous approaches did not consider those issues and performed instead global filtering techniques that broadly miss the rich pattern structure.

² The more complex multiparametric general/specific transformations are studied in section 3.



© Olivier Lartillot.

Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). **Attribution:** Olivier Lartillot. “IN-DEPTH MOTIVIC ANALYSIS BASED ON MULTIPARAMETRIC CLOSED PATTERN AND CYCLIC SEQUENCE MINING”, 15th International Society for Music Information Retrieval Conference, 2014.




Figure 1. Patterns found in a sequence of symbols. Below the sequence, each row represents a different pattern class with the occurrences aligned to the sequence. Thick black lines correspond to closed patterns (the upper one is the maximal pattern), grey lines to prefixes of closed patterns, and thin lines to non-closed patterns.




Figure 2. Closed patterns found in a cyclic sequence of symbols. The occurrences of the pattern shown in thick lines do not overlap, whereas those shown in thin lines do.

1.3 New approach

In this paper, we propose a simplified description and modelling of this exhaustive pattern mining approach. In section 2, we present the problem of closed pattern mining on the simple case of monoparametric string analysis, introduce a simplified algorithmic implementation, and present a new way to simply justify the interest of the approach. In section 3, the approach is generalised to the multidimensionality of the musical parametric space. Section 4 discusses pattern cyclicity and presents a new simple model that solves this issue. In section 5, the interest of the method is shown through the analysis of a piece of music from the JKUPDD database.

2. CORE PRINCIPLES OF THE MODEL

2.1 Advantages of incremental one-pass approach

As explained in the previous section, testing the closedness of a pattern requires comparing its number of occurrences with those of all the more specific patterns. Previous computer science researches in closed pattern mining (one recent being [9]) incrementally construct the closed patterns dictionary while considering the whole document to be analysed (in our case, the piece of music). This requires the design of complex algorithms to estimate the number of occurrences of each possible pattern candidate.

We introduced in [6] a simpler approach based on an incremental single pass throughout the document (i.e., from the beginning to the end of the piece of music), during which the closed pattern dictionary is incrementally constructed: for each successive note n in the sequence, all

patterns in the subsequence ending to that note n are exhaustively searched for. The main advantage of the incremental approach is based on the following property.

Lemma 2.1 (Closed pattern characterisation). *When following the incremental approach, for any closed pattern P , there exists a particular moment in the piece of music where an occurrence O of P can be inferred while no occurrence of any more specific pattern can be inferred.*

Proof. There are three alternative conditions concerning the patterns more specific than P :

- There is no pattern more specific than P . In this case, the observation is evident.
- There is only one pattern S more specific than P . For instance, in Figure 3, $S = ABCD$ is more specific than $P = CD$. Since P is closed, it has more occurrences than S , so there exists an occurrence of P that is not occurrence of S .
- There are several patterns S_1, \dots, S_n more specific than P . For instance, in Figure 1, $S_1 = ABCDE$ and $S_2 = ABCDECDE$ are both more specific than $P = CDE$. As soon as two different more specific patterns S_1 (one or several time) and S_2 (first time) have appeared in the sequence, pattern P can be detected, since it is repeated in S_1 and S_2 , but S_2 is not detected yet, since it has not been repeated yet.

□

As soon as we detect a new pattern repetition, such that for that particular occurrence where the repetition is detected, there is no more specific pattern repetition, we can be sure that the discovered pattern is closed.

When considering a given pattern candidate at a given point in the piece of music, we need to be already informed about the eventual existence of more specific pattern occurrences at the same place. Hence, for a given note, patterns need to be extended *in decreasing order of specificity*.

To details further the approach, let's consider in a first simple case the *monoparametric contiguous string* case, where the main document is a sequence of symbols, and where pattern occurrences are made of contiguous substrings. In this case, 'more general than' simple means 'is a subsequence of'. In other words, a more general pattern is a *prefix* or/of a *suffix* of a more specific pattern. Let's consider these two aspects separately:

- Since the approach is incremental, patterns are constructed by incrementally extending their *prefixes* (in grey in Figure 1). Patterns are therefore represented as chains of prefixes, and the pattern dictionary is represented as a prefix tree. In this paradigm, if a given pattern P is a prefix of a closed pattern S , and if both have same number of occurrences, the prefix P can still be considered as a closed pattern, in the sense that it is an intermediary state to the constitution of the closed pattern S .



CDABCDEABCDE

Figure 3. Closed patterns found in a sequence of symbols. The occurrence during which a pattern is discovered is shown in black. Dashed extensions indicate two possible pattern extensions when integrating the last note.

- The closedness of a pattern depends hence solely on the patterns to which it is a *suffix*. Thanks to the incremental one-pass approach, these more specific patterns are already inferred. The only constraint to be added is that when a given note is considered, the candidate patterns should be considered in decreasing order of specificity, i.e. from the longest to the shortest (which are suffixes of the longer ones). For instance, in Figure 3, when analysing the last note, E, there are two candidate patterns for extension, ABCD and CD. Since we first extend the most specific pattern ABCDE, when considering then the more general pattern CD, extension CDE is found as non-closed and thus not inferred.

2.2 Algorithmic details

Following these principles, the main routine of the algorithm simply scans the musical sequence chronologically, from the first to the last note. Integrating a new note consists in checking:

- whether pattern occurrence(s) ending at the previous note can be extended with the new note,
- whether the new note initiates the start of a new pattern occurrence.

The extension of a pattern occurrence results from two alternative mechanisms:

Recognition the new note is recognised as a known extension of the pattern.

Discovery the new note continues the occurrence in the same way that a previous note continued an older occurrence of the pattern: the pattern is extended with this new common description, and the two occurrences are extended as well.

Concerning the discovery mechanism, the identification of new notes continuing older contexts can be implemented using a simple associative array, storing the note following each occurrence according to its description. This will be called a *continuation memory*. Before actually extending the pattern, we should make sure that the extended pattern is closed.

2.3 Specific Pattern Class

Searching for all closed patterns in a sequence, instead of all possible patterns, enables an exhaustive pattern analysis without combinatorial explosion: all non-closed patterns

can be deduced from the closed pattern analysis. Yet, the set of closed patterns can remain quite large and the exhaustive collection of their occurrences can become cumbersome. [6] proposes to limit the analysis, without any loss of information, to closed patterns' *specific classes*, which correspond to pattern occurrences that are not included in occurrences of more specific patterns. For instance, in Figure 3, the specific class of CD contains only its first occurrence, because the two other ones are superposed to occurrences of the more specific pattern ABCDE.

We propose a simpler model for the determination of specific class of closed patterns. Non-specific occurrences are regenerated whenever necessary. Because occurrences of a given pattern are not all represented, the notes following these occurrences are not memorised, although they could generate new pattern extensions. To circumvent this issue, the extension memory related to any given pattern contains the extensions not only of that pattern but also of any more specific pattern.

3. MULTIPARAMETRIC PATTERN MINING

The model presented in the previous section searches for sequential patterns on monoparametric sequences, composed of a succession of symbols taken from a given alphabet. Music cannot be reduced to unidimensional parametric description.

3.1 Parametric space

The problem needs to be generalised by taking into account three main aspects:

- Notes are defined by a hierarchically structured combination of parameters (diatonic and chromatic pitch and pitch class, metrical position, etc.).
- Notes are defined not only in terms of their absolute position on fixed scales, but also relatively to a given local context, and in particular with respect to the previous notes (defining pitch interval, gross contour, rhythmic values, etc.). These interval representations are also hierarchically structured. Gross contour, for instance, is a simple description of the inter-pitch interval between successive notes as “increasing”, “decreasing” or “unison”. Matching along gross contour enables to track intervallic augmentation and diminution. For instance, in the example in section 5, the first interval of the fugue subject is either a decreasing third or a decreasing second. The actual diatonic pitch interval representation differs, but the gross contour remains constantly “decreasing”.
- A large part of melodic transformations can be understood as repetitions of sequential patterns that do not follow strictly all the parametric descriptions, but only a subset. For instance, a rhythmical variation of a melodic motif consists in repeating the pitch sequence, while developing the rhythmical part more freely.

[6] proposes to integrate both absolute note position and relative note interval into a single parametric space. This enables to define a motive and any occurrence as a simple succession of parametric descriptions. [6] also shows the importance of heterogeneous patterns, which are made of a succession of parameters that can each be defined on different parametric dimensions. For instance, the subject of the fugue analysed in section 5 is heterogeneous, as it starts with a gross contour interval followed by more specific descriptions. In the multiparametric paradigm, a pattern G is more general than a pattern S if it is a suffix of S and/or the successive parametric descriptions of the patterns are equal or more general than the related parametric descriptions in pattern P .

3.2 Motivic/thematic class as “paradigmatic sheaf”

Extending the exhaustive method developed in the previous section to this heterogeneous pattern paradigm enables to describe all possible sequential repetitions along all parametric dimensions. This leads to very detailed pattern characterisation, describing in details the common sequential descriptions between any pair of similar motif. However, a more synthetic analysis requires structuring the set of discovered patterns into motivic or thematic classes. Manual motivic taxonomy of these discovered patterns has been shown in [7].

We have conceived a method for the collection of all patterns belonging to a same motivic or thematic class. Starting from one pattern seed, the method collects all other patterns that can be partially aligned to the seed, as well as those that can be aligned to any pattern thus collected. Patterns are searched along the following transformations:

- More general patterns of same length
- More specific patterns: only the suffix that have same length that the pattern seed is selected.
- Prefixes of pattern seed can be used as pattern seeds too: they might contain additional sets of more general and more specific patterns of interest.
- Pattern extensions, leading to a forking of the motivic or thematic class into several possible continuations

All the patterns contained in the bundle remain informative in the way they show particular commonalities between subset of the motivic/thematic class, as shown in the analysis in section 5.

3.3 Heterogeneous pattern mining

A parametric description of a given note in the musical sequence instantiates values to all fields in the parametric space. Values in the more general fields are automatically computed from their more specific fields. A parametric description of a note in a pattern instantiates values to some fields in the space, the other indeterminate fields corresponding to undefined parameters. Values can be assigned to more general fields, even if no value is assigned

to their corresponding more specific fields. Methods have been implemented that enable to compare two parametric descriptions, in order to see if they are equal, or if one is subsumed into the other, and if not, to compute the intersection of the two descriptions.

The multiparametric description is integrated in the two core mechanisms of the incremental pattern mining model as follows:

Recognition As before, the observed parametric description of the new note is compared to the descriptions of the patterns’ extensions. If the pattern extension’s description fits only partially, a new more general pattern extension is created (if not existing yet) related to the common description.

Discovery The continuation memory is structured in the same way as the parametric space: for each possible parametric field, an associative memory stores pattern continuations according to their values along that particular parametric field. As soon as a stored pattern continuation is identified with the current note along a particular parametric field, the complete parametric description common to these two contexts is computed, and the pattern extension is attempted along that common parametric description. As before, a pattern is extended only if the extended pattern is closed.

4. PATTERN CYCLICITY

A solution to the problem of cyclicity introduced in section 1.2 was proposed in [6] through the formalisation of *cyclic patterns*, where the last state of the chain representing the pattern is connected back to its first state, formalising this compelling expectation of the return of the periodic pattern. One limitation of the approach is that it required the explicit construction of cyclic pattern, which demanded contrived algorithmic formalisations. The problem gets even more difficult when dealing with multiparametric space, in particular when the pattern is only partially extended, i.e., when the expected parametric description is replaced by a less specific parametric matching, such as in the musical example shown in Figure 4. In this case, a more general pattern cyclic needs to be constructed, leading to the inference of a complex network of pattern cycles particularly difficult to conceptualise and implement.

We propose a simpler approach: instead of formalising *cyclic patterns*, pattern cyclicity is represented on the pattern *occurrences* directly. Once a successive repetition of a pattern has been detected, such as the 3-note pattern starting the musical example in Figure 4, the two occurrences are fused into one single chain of notes, and all the subsequent notes in the cyclic sequence are progressively added to that chain. This *cyclic chain* is first used to track the development of the new cycle (i.e., the third cycle, since there were already two cycles). The tracking of each new cycle is guided by a model describing the expected sequence of musical parameters. Initially, for the third cycle, this




Figure 4. Two successive repetitions of a pattern, at the beginning of the musical sequence, characterised by a pitch sequence (G, C, Eb, and back to G), a pitch interval sequence (ascending perfect fourth (+3), ascending minor third (+2m) and descending minor sixth (-5m)), and a rhythmical sequence made of a succession of 8th notes. This successive repetition leads to the inference of a cyclic chain, indicated at the bottom of the figure. When this cycle is initially inferred, at note 7, the model of the cycle, represented above “cycle 3”, corresponds to the initial pattern description. At note 10, some descriptions expected by the model (indicated in bold italics) are not fulfilled, but a more general description is inferred (descending gross contour (-)). Consequently, the next cycle (4)’s model is generalised accordingly. At note 13, a new regularity is detected, due to the repetition of pitch Ab and of descending perfect fifth (-4). Consequently, the next cycle (5)’s model is specialised accordingly.

model corresponds to the pattern that was repeated twice in the two first cycles.

- If the new cycle scrupulously follows the model, this same model will be used to guide the development of the subsequent cycle.
- If the new cycle partially follows the model (such as the modification, at the beginning of bar 2 in Figure 4, of the decreasing sixth interval, replaced by a more general decreasing contour), the model is updated accordingly by replacing the parameters that have not been matched with more general parameters.
- If the new cycle shows any new pattern identification with the previous cycle (such as the repetition of pitch Ab at the beginning of cycles 4 and 5 in Figure 4), the corresponding descriptions are added to the model.
- If at some point, the new note does not match at all the corresponding description in the model, the cyclic sequence is terminated.

This simple method enables to track the cyclic development of repeated patterns, while avoiding the combinatorial explosion inherent to this structural configuration.

5. TESTS

The model described in this paper is applied to the analysis of the Johannes Kepler University Patterns Develop-

ment Database (JKUPDD-Aug2013), which is the training set part of the MIREX task on Discovery of Repeated Themes & Sections initiated in 2013, and made publicly available, both symbolic representation of the scores and ground-truth musicological analyses [4].

This section details the analysis of one particular piece of music included in the JKUPDD, the 20th Fugue in the Second Book of Johann Sebastian Bach’s *Well-Tempered Clavier*. The ground truth consists of the two first bars of the third entry in the exposition part along the three voices that constitute this fugue [1]. The third entry is chosen because it is the first entry where the subject and the two countersubjects are exposed altogether. To each of these three ground-truth patterns (the subject and the two countersubjects in this two-bar entry), the ground-truth data specifies a list of occurrences in the score.

Figure 5 shows the thematic class related to ground-truth pattern #1, i.e., the fugue’s subject. This is detected by the model as one single motivic/thematic class, i.e., one complete paradigmatic sheaf, resulting from the bundling method presented in section 3.2. All occurrences indicated in the ground truth are retrieved. The patterns forming this thematic class are longer than the two-bar motif indicated in the ground truth. The limitation of all subjects and counter-subjects in the musicological analysis to two bars stems from a theoretical understanding of fugue structure that cannot be automatically inferred from a direct analysis of the score.

The analysis offered by the computational model offers much richer information than simply listing the occurrences of the subjects and countersubjects. It shows what musical descriptions characterise them, and details particular commonalities shared by occurrences of these subjects and countersubjects. For instance entries M1 and U1 belong to a same more specific pattern that describes their particular development. L1, U1 and U3 start all with a decreasing third interval, and so on.

The model presented in this paper does not yet integrate mechanisms for the reduction of ornamentation, as discussed in the next section. The only melodic ornamentation appearing in pattern #1 is the addition of a passing note after the first note of occurrences L2 and L3. This leads to a small error in the model’s results, where the first actual note is not detected.

The thematic class related to ground-truth pattern #2, which is the first countersubject, is extracted in the same way, forming a paradigmatic sheaf. The pattern class given by the model corresponds mostly to the ground truth. Here again, some occurrences present similar extensions that are inventoried by the model, although they are ignored in the ground truth. The last occurrence, which is a suffix of the pattern, is also detected accordingly. On the other hand, the second last occurrence is not properly detected, once again due to the addition of passing notes.

Pattern #3, which is the second countersubject, is more problematic, because it is only 7 notes long. Several other longer patterns are found by the model, and the specificity of pattern #3 is not grounded on characteristics purely re-

Figure 5. Entries of the subject in Bach’s Fugue, as found by the model. The fugue has three voices: upper (U), middle (M) and lower (L). In each entry is slurred the part actually indicated in the ground-truth description of the subject. The model proposes a longer description of the subject, that is particularly developed in M1 and U1.

lated to pattern repetition. As aforementioned, the ground-truth selection of these three patterns are based on principles related to fugue rules, namely the synchronised iteration of the three patterns along the separate voices. It seems questionable to expect a general pattern mining algorithm non-specialised to a particular type of music to be able to infer this type of configuration.

6. CONCLUSION

The approach is incremental, progressively analysing the musical sequence through one single pass. This enables to control the structural complexity in a way similar to the way listeners perceive music.

Gross contour needs to be constrained by factors related to local saliency and short-term memory. The integration of more complex melodic transformation such as ornamentation and reduction is currently under investigation. Motivic repetition with local ornamentation is detected by reconstructing, on top of “surface-level” monodic voices, longer-term relations between non-adjacent notes related to deeper structures, and by tracking motives on the resulting syntagmatic network. More generally, the analysis of

polyphony is under study, as well as the application of the pattern mining approach to metrical analysis. The system, implemented in Matlab, is made publicly available as part of *MiningSuite*³, a new open-source framework for audio and music analysis.

7. ACKNOWLEDGMENTS

This work was funded by an Academy of Finland research fellowship at the Finnish Centre of Excellence in Interdisciplinary Music Research at the University of Jyväskylä. The research is continued in the context of the European project *Learning to Create (Lrn2Cre8)*, which acknowledges the financial support of the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET grant number 610859.

8. REFERENCES

- [1] S. Bruhn. *J.S. Bach’s Well-Tempered Clavier: in-depth analysis and interpretation*, Mainer International, Hong Kong, 1993.
- [2] E. Cambouropoulos. *Towards a General Computational Theory of Musical Structure*, PhD thesis, University of Edinburgh, 1998.
- [3] E. Cambouropoulos. “Musical parallelism and melodic segmentation: A computational approach,” *Music Perception*, 23(3), pp. 249–268, 2006.
- [4] T. Collins. MIREX 2013: Discovery of Repeated Themes and Sections, 2013. http://www.music-ir.org/mirex/wiki/2013:Discovery_of_Repeated_Themes_&_Sections Accessed on 14 August 2014.
- [5] D. Conklin, and C. Anagnostopoulou. “Representation and Discovery of Multiple Viewpoint Patterns,” *Proceedings of the International Computer Music Conference*, 2001.
- [6] O. Lartillot. “Efficient Extraction of Closed Motivic Patterns in Multi-Dimensional Symbolic Representations of Music,” *Proceedings of the International Symposium on Music Information Retrieval*, 2005.
- [7] O. Lartillot. “Taxonomic categorisation of motivic patterns,” *Musicae Scientiae*, Discussion Forum 4B, pp. 25–46, 2009.
- [8] D. Meredith, K., Lemström, and G. Wiggins. “Algorithms for discovering repeated patterns in multidimensional representations of polyphonic music,” *Journal of New Music Research*, 31(4), pp. 321–345, 2002.
- [9] J. Wang, J. Han, and C. Li. “Frequent closed sequence mining without candidate maintenance,” *IEEE Transactions on Knowledge and Data Engineering*, 19:8, pp. 1042–1056, 2007.

³ Available at <http://code.google.com/p/miningsuite/>.