



Royal Netherlands Academy of Arts and Sciences (KNAW) KONINKLIJKE NEDERLANDSE AKADEMIE VAN WETENSCHAPPEN

Discovering repeated patterns in music: potentials, challenges, open questions

Janssen, B.; de Haas, W.B.; Volk, A.; van Kranenburg, P.

published in

Proceedings of the 10th International Symposium on Computer Music Multidisciplinary Research
2013

document version

Peer reviewed version

[Link to publication in KNAW Research Portal](#)

citation for published version (APA)

Janssen, B., de Haas, W. B., Volk, A., & van Kranenburg, P. (2013). Discovering repeated patterns in music: potentials, challenges, open questions. In *Proceedings of the 10th International Symposium on Computer Music Multidisciplinary Research* (pp. 225-240). Laboratoire de Mecanique et d'Acoustique.

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 KNAW 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 KNAW public portal.

Take down policy

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

E-mail address:

pure@knaw.nl

Discovering repeated patterns in music: state of knowledge, challenges, perspectives

Berit Janssen^{1,2}, W. Bas de Haas³, Anja Volk³, and Peter van Kranenburg¹

¹ Meertens Institute, Amsterdam, the Netherlands

² Institute for Language, Logic and Computation, Amsterdam University, the Netherlands

³ Department of Information and Computing Sciences, Utrecht University, the Netherlands

`berit.janssen@meertens.knaw.nl`

Abstract. This paper discusses the current state of knowledge on musical pattern discovery. Various studies propose computational methods to find repeated musical patterns. Our detailed review of these studies reveals important challenges in musical pattern discovery research: different methods have not yet been directly compared, and the influence of music representation and filtering on the results has not been assessed. Moreover, we need a thorough understanding of musical patterns as perceived by human listeners. A sound evaluation methodology is still lacking. Consequently, we suggest perspectives for musical pattern discovery: future research can provide a comparison of different methods, and an assessment of different music representations and filtering criteria. A combination of quantitative and qualitative methods can overcome the lacking evaluation methodology. Musical patterns discovered by human listeners form a reference, but also an object of study, as computational methods can help us understand the criteria underlying human notions of musical repetition.

Keywords: musical pattern discovery; music analysis; music cognition; music information retrieval

1 Introduction

Repetitions are a fundamental structuring principle in many musical styles. They guide the listener in their experience of a musical piece through creating listening experiences, and facilitate the recall process [20, p.228 ff.]. As such, the study of repetition is an important research topic in many fields of music research, and computational methods enable researchers to study musical repetitions quantitatively in large music collections.

Musical pattern discovery is important in several areas of music research. In Music Information Retrieval, repetitions have been used as indicators of musical segmentation, or to find themes or choruses in large databases [37]. In Music Analysis, analytical approaches based on repetition, for instance Réti's motivic

analysis [40], have been formalized and evaluated by developing a computer model [3]. In Folk Music Research, computational discovery of shared patterns between folk song variants offers the potential to detect moments of stability, i.e. melodic elements that change relatively little through the process of oral transmission [45]. Hypotheses on memory, recall and transmission of melodies can be tested on large databases once stable melodic patterns can be discovered automatically [21].

As we will show in this article, there are many different kinds of repetition that different researchers investigate: large repeated structures, such as themes, chorusses, or stanzas; smaller repeated units, such as motifs; but also building blocks of improvised music, such as formulae or licks. For these different purposes, and for different genres, the authors of the discussed studies have formalised repetition in different ways. What may be considered as a variation or as musically unrelated depends on a great number of factors, factors which yet need to be understood [44].

A computational method to discover musical repetitions can contribute to an understanding of principles of repetition and variation. Using computational methods, researchers can model and test knowledge on the cognitive principles underlying musical repetition. Cognitive and computational models can cross-pollinate each other in such an overarching research question, as argued, for instance, by Honing [18, 19]. In this paper, we explore computational methods for the discovery of repeated musical patterns, which we will refer to as musical pattern discovery.

Currently, the knowledge of musical pattern discovery is dispersed across different fields of music research. Miscellaneous studies present various approaches to the problem, but there is no systematic comparison of the proposed methods yet. Moreover, the influence of music representation, and filtering of algorithmic results on the success of musical pattern discovery is currently unknown. This lack of comparative assessment is further complicated by the lack of a sound methodology for the evaluation of musical pattern discovery.

This paper provides a comprehensive overview, review and discussion of the field of musical pattern discovery. We present the essence of assorted studies, and we proceed to clarify the relationships between different methods, proposing a taxonomy of musical pattern discovery approaches, and discussing various studies according to the criteria of music representation, filtering, reference data, and evaluation. Furthermore, we identify current challenges of musical pattern discovery and conclude steps to overcome these challenges.

2 State of knowledge on musical pattern discovery

We conducted a comprehensive literature survey on musical pattern discovery. In this section, we provide an overview of the various studies. First, we discuss the methods that the surveyed studies used, introducing a taxonomy within which the different approaches can be placed. We continue by considering different forms of music representation, different methods of filtering algorithmic results,

and close with an overview of reference data and evaluation methods used in the studies.

Our focus in the survey was on studies using symbolic music representations. Several studies work with audio representations [7, 33], but we will not discuss the problems related to transforming audio representations for musical pattern discovery. For this, and other methods for the audio domain, we refer the reader to the overview by Klapuri [23]. Moreover, our focus is on pattern discovery rather than pattern matching; pattern matching refers to locating predefined patterns in a large document or corpus [41]. Some applications of pattern matching are locating fugue subjects and counter-subjects in Bach fugues [16], or matching musical pieces against musical keys [30]. These are different problems than the identification of motifs, themes, and other musical structures without prior knowledge of possible candidates.




Fig. 1. A schematic representation of a taxonomy for musical pattern discovery methods. For each category, relevant studies are listed.

We present our overview of studies on musical pattern discovery in Tables 1 and 2. The table’s columns refer to the study, naming the first author and year (please refer to the references for the full list of authors). The second column mentions the goals of the authors in exploring musical pattern discovery. We list the musical pieces that each study used. Further categories, which we will discuss in more detail below, are the method the studies employed, the music representation that was used, what kind of filtering of the algorithmic results

was performed, the reference of musical patterns against which the results were compared, represented by first author and year, and finally, which evaluation methods were applied. Empty cells denote that the category in question is not applicable to a particular study.

2.1 Methods

As can be gleaned from the various goals of the studies presented in Tables 1 and 2, the interest for musical pattern discovery is diverse, spanning different music research disciplines and musical genres. This might be the reason why various authors present their algorithms without stating how their method relates to some of the other research on musical pattern discovery.

To assess the relationship of different methods, we propose a taxonomy in which the various methods can be placed in perspective, which is represented in Figure 1.

String-based or geometric methods A piece of music can be represented by a series - or string - of tokens. For instance, the notes of a melody could be represented by a string of pitches (A; G; A; D), as MIDI note numbers (57; 55; 57; 50), or as tokens representing both pitch and duration of the note ((A,1.5),(G,0.5),(A,1.0),(G,1.0)). These multiple possibilities will be further discussed in the section on music representation below.

One approach to musical pattern discovery is to search for identical subsequences of tokens in a string representation of a melody or multiple melodies. This approach has been derived from techniques developed within Computational Biology to compare gene sequences. Gusfield [17] provides a thorough overview of these techniques. Most of the studies presented here rely on such algorithms. We will refer to these approaches as string-based methods.

Meredith's *Structure Induction Algorithms* (SIA) [34] present an alternative to these algorithms, using a geometric approach. In geometric methods, a melody is considered as a shape in an n-dimensional space. Repeated patterns are then identified as (near-)identical shapes. According to Meredith, geometric methods handle polyphonic pieces more elegantly than string-based methods [34, p.328].

Next to Collins et al. [10], who introduced improvements of Meredith's Structure Induction Algorithms, we also place the studies by Buteau and Viperman [4] and Szeto and Wong [43] in the category of geometric approaches, as both studies also use a notion of musical patterns as shapes. Szeto and Wong represent these geometric relationships as graphs [43].

Use of indexing structures Within the string-based methods we can define two categories: namely, whether or not indexing structures are used. Such indexing structures typically are a graph of tree-like structure, in which the repetitions contained in a string are represented, and can therefore be efficiently found. We will first present those studies that do not use indexing structures for the music pattern discovery, then those which do use them.

Study	Goals	Musical pieces	Method	Music representation	Filtering	Reference data	Evaluation
Buteau2009 [4]	Computer-aided motivic analysis	Schumann's <i>Träumerei</i>	topological motifs as shapes in n-dimensional space	pitch and onset	shape clustering	Repp1992 [39]	Qualitative comparison with music theorist Repp's segmentation of Schumann's <i>Träumerei</i>
Cambourpoulos2006 [5]	Segmentation based on musical repetitions	Sonatina No. 2 in C Major by Anton Diabelli (soprano voice)	set partitioning method [14]	pitch intervals, step-leap	frequency of occurrence, pattern length	Konari2001 [27]	Quantitative comparison with human segmentations of a sonatina by Anton Diabelli
Cambourpoulos2007 [6]	Finding motifs in musical pieces		Searching maximal pairs in suffix trees	pitch interval, step-leap representation of pitch			
Chou1996 [7]	Content-based retrieval of music from large data bases		Using a form of radix trees ("PAT trees") as indexing structure	representation of chords			
Collins2010 [10]	Finding translational patterns in musical pieces to support music analysis	two sonatas by Scarlatti, two Bach preludes (polyphonic)	Structure from Algorithms (extension of Meredith2002)	Inductive-pitch interval, on-pitch interval, on-pitch set	"compactness crawling": finding patterns with great number of datapoints within pattern classes	Collins2010 [10]	Quantitative comparison with motif annotations in two sonatas by Scarlatti and two Bach preludes; using Precision and Recall
Conklin2011 [13]	Classifying Cretan folk songs based on characteristic patterns	106 mono-phonetic Cretan folksongs	sequential mining: maximally general distinctive pattern (MGDP) in a search tree	pitch intervals	frequency of occurrence		Qualitative analysis of patterns as parts of Cretan folk music genres
Karydis2006 [22]	Finding themes in large music databases	Assorted MIDI files (unspecified)	Finding maximum-length repeating patterns with set partitioning method	pitch			Computation speed
Knopke2009 [25]	Analysis of Palestrina's contrapuntal structure	masses by Palestrina (separated into voices and phrases)	4 suffix arrays for each phrase (indicated by rests), mapped to UTF-8 tokens, retrograde, retrograde inversion	pitch and duration intervals, mapped to UTF-8 tokens			Qualitative analysis of selected patterns, especially longest patterns and most frequent patterns

Table 1. Overview of research on musical pattern discovery in music

Study	Goals	Musical pieces	Method	Music representation	Filtering	Reference data	Evaluation
Lartillot2005 [28]	Assisting music analysis with computational discovery of motifs	monophonic melodies from different genres	construction of a pattern tree, which can have cyclic structures	pitch intervals, duration	maximally specific filtered patterns, during the tree construction		Qualitative analysis of selected patterns, discovered in monophonic melodies from different genres
Lee1999 [29]	Content-based retrieval of music from large data bases	100 songs in MIDI format (unspecified)	Combined and independent suffix trees for pitch interval and duration, introducing Twin Suffix Trees and Grid-Twin Suffix Trees	pitch interval, quantized duration			Computation speed
Meek2001 [33]	Discovering themes in Western art music	60 polyphonic pieces from different epochs of Western art music	Stream segregation; matching keys representing pitch intervals in a base 26 numerical system	pitch intervals	frequency of occurrence; hill-climbing optimisation of different pattern features	Barlow1948 [1]	Quantitative comparison of selected examples against <i>Dictionary of musical themes</i> (Barlow): the evaluation procedure is not described
Meredith2002 [34]	Assisting music analysts and composers by finding repeated patterns in music	52 polyphonic pieces, mostly from Bach's <i>Wohltun-pervertes Clavier</i>	Structure Induction Algorithms: tones represented as points in n-dimensional space	pitch intervals, onset	Suggestions, not reported		Demonstration of one discovered pattern
Nieto2012 [35]	Finding motifs in musical works	Monophonic melodies from 6 musical pieces (Bach, Mozart, Beethoven, Haydn)	string-based pattern discovery; Gestalt rules applied for candidate motifs	pitch intervals, duration	clustering candidate motifs		Score based on amount of overlap between annotated motifs in the pieces and those discovered by algorithm
Rolland1999 [41]	Automatizing discovery in jazz improvisation	10 Charlie Parker solos	Knuth-Morris-Pratt approximate matching with star center algorithm, alignment for similarity	duration, intervals, metric levels, contour	Frequency of occurrence	Owens1974 [36]	Qualitative comparison of selected patterns with Owens' catalogue of Charlie Parker's patterns
Szeto2006 [43]	Discovering structures in post-tonal music	Arnold Schoenberg's Klavierstücke, op. 11 no.1	Graph with notes as vertices, edges linking simultaneous and sequential events	pitch class, onset	Clustering of graph vertices	Wittlich1974 [46]	Quantitative comparison with musicological analysis (Wittlich1974): counting number of matches and approximate matches

Table 2. Overview of research on musical pattern discovery in music (continued)

The simplest string-based approach to finding repeated patterns in a melody M consists of sliding all possible patterns P past M , and recording all found matches. This approach is taken by Nieto and Farbod [35]. There are some extensions of this simple approach, which skip some comparison steps between P and M without missing any relevant patterns. One of these extensions, the algorithm by Knuth, Morris and Pratt [26], has been applied by Rolland [41] for musical pattern discovery.

Yet another approach is Crochemore’s set partitioning method [14], which recursively splits the melody M into sets of repeating tokens. Cambouropoulos [5] used this method to find maximally repeating patterns (i.e. repeated patterns which cannot be extended left or right and still be identical) in musical pieces. Karydis et al. [22] refine the set-partitioning approach to find only the longest patterns for each musical piece, with the intuition that these correspond most closely to musical themes.

Meek and Birmingham’s [33] algorithm transforms all possible patterns up to a maximal pattern length to keys in a radix 26 system (representing 12 intervals up or down, unison, and a 0 for the end of the string). After a series of transformations, which consolidate shorter into longer patterns; identical patterns are encoded by the same numerical keys.

There are number of studies which do use indexing structures [6, 7, 13, 25, 28, 29]. Knopke and Jürgensen [25] use suffix arrays representing phrases of Palestina masses. Conklin and Anagnostopoulou [13] use a sequence tree to represent search spaces of patterns in Cretan folk songs, which is pruned based on the patterns’ frequency of occurrence. Lartillot [28] employs a pattern tree to model musical pieces; what is special about this indexing structure is that it allows for cyclic structures within the graph, which can capture repetitions of short patterns, forming building blocks within larger repeated structures.

Exact or approximate matching Next to searching for exact matches, approximate matching is also of great interest to musical pattern discovery. Rhythmic, melodic and many other conceivable variations are likely to occur, such as the insertion of ornamentations during a repetition, the speeding up or slowing down of a musical sequence, deviations in pitch, or transpositions.

Many different ways to define approximate matching have been proposed in theory, see [8] for an overview. To our knowledge, only Rolland [41] applied approximate matches in musical pattern discovery, using the Levenshtein edit distance [31] to compare a pattern with a match candidate. The candidate is then accepted as a match if the edit distance remains under a user-defined threshold. Implicitly, however, approximate matching can also be achieved through more abstract music representations, as Cambouropoulos et al. [6] point out. We will discuss this in more detail below.

2.2 Music representation

There are different musical dimensions to be considered for comparisons of musical patterns: rhythm, pitch, but also dynamics, timbre, and many more. More-

over, there are different conceivable abstraction levels at which to represent these dimensions: in terms of absolute values; in terms of categories or classes; in terms of contours indicating only the direction of change; among others. Compare also to Conklin’s [11] notion of musical viewpoints.

A glance at the music representation column reveals that the majority of the presented studies on musical pattern discovery use pitch or pitch intervals as the music representation, in some of the studies this is combined with rhythmic representations such as note onset or duration.

Chou et al. [7] derive a more abstract representation from the musical surface: they represent polyphonic musical pieces as chords, which are automatically extracted from simultaneous pitches, even though the paper does not report the results of such an approach.

Three studies [41, 6, 29] suggest multiple music representations. Rolland [41] allows the users of his *FLEXPAT* software to switch between different music representations, but he does not report how this influences the results of his musical pattern discovery algorithm. Cambouropoulos et al. [6] suggest to compare a pitch interval representation with a more abstract step-leap representation, but results of these two representations are not discussed by the authors. Lee and Chen [29] state that pitch and duration can be either represented in independent suffix trees, or in a combined suffix tree which contains tuples of pitch and duration values. They do not find either approach satisfying, so they suggest two new indexing structures, *Twin Suffix Trees*, and *Grid-Twin Suffix Trees*, as possible alternatives. They do not report any results produced by these different representations.

2.3 Filtering

A frequently described problem in musical pattern discovery is the great amount of algorithmically discovered pattern as compared to the patterns that would be considered relevant by a human analyst. For the task of computer-aided motivic analysis, Marsden recently observed that “... the mathematical and computational approaches find many more motives and many more relationships between fragments than traditional motivic analysis.” [32]

Therefore, most of the presented studies employ a filtering step, which is supposed to separate the wheat from the chaff. One approach is to consider only relatively long patterns. Cambouropoulos [5] filters according to pattern lengths. Likewise, Collins’ *compactness trawling* to refine the result of Structure Induction Algorithms [10] is based on such a filtering.

Another frequently employed filtering method is based on the assumption that patterns which occur more often might also be considered as more important by human analysts, and is applied in several studies [5, 41].

Conklin and Anagnostopoulou [13] are also interested in a pattern’s frequency of occurrence, but weigh it against its frequency in a collection of contrasting music pieces, the *anticorpus*. This process is designed to favour patterns that are characteristic for the analysed piece. Conversely, also patterns which are

characteristically not represented in specific pieces or genres can be interesting for music researchers [12].

Nieto and Farbod filter according to Gestalt rules during the search process, which means that patterns containing relatively long notes or rests, or relatively large intervals will be rejected as candidates [35].

Lartillot’s detection of maximally specific patterns [28] filters for patterns which repeat in different musical domains: a pattern which repeats the rhythm as well as the pitch of another pattern is considered as more specific than one that repeats in one of these domains alone.

It is also possible to employ several filtering parameters, and adjust their relative weights using an optimization algorithm. Meek and Birmingham [33] take this approach. They let the algorithm select those filtering parameters which give the best agreement between the discovered patterns and Barlow and Morgenstern’s *Dictionary of Musical Themes* [1].

2.4 Reference data

Some of the presented studies use annotated musical patterns to evaluate the results of their musical pattern discovery algorithm, which we will denote as reference data. Such reference data ranges from overviews of frequently used licks in jazz improvisation [36] to themes in Western art music [1]. This reference data is typically assembled by domain specialists, annotating what they consider the most relevant patterns of the analysed music collection.

Such reference data can serve to evaluate musical pattern discovery: if there is a good agreement between an algorithm and a reference, the algorithm emulates human judgements on relevant patterns, and might therefore be more useful for tasks such as assisting music analysis for this genre.

2.5 Evaluation

Most studies discussed here evaluate qualitatively: the discovered patterns are scrutinised. Typically, selected examples are presented to the reader in this case. If there is no reference against which the results can be compared, the researchers’ and readers’ judgements form the post hoc reference.

For available reference data, several researchers evaluate the discovered patterns quantitatively by counting the number of agreements between algorithm and reference data [10, 33, 35, 43].

Sometimes, the reference data takes the form of a segmentation. In this case, the starting and end positions of algorithmically discovered patterns can be compared with segmentation boundaries [4, 5].

Several quantitative evaluation measures have been suggested for musical pattern discovery, but there seems to be no common methodology between the reported studies. This is one of the points to which we will turn in the next sections, in which we will discuss the current challenges in the field of musical pattern discovery.

3 Challenges of musical pattern discovery

We have now observed that in various studies different approaches to musical pattern discovery have been proposed. There are some common challenges which can be derived from these observations, which will be discussed below, following the same order as the previous section.

3.1 Methods

The approaches to musical pattern discovery presented above all aim to find specific kinds of patterns, in specific genres and for specific applications. This is a good starting point, as a restriction to one part of the research field enables researchers to formulate more concrete questions, and to interpret results more easily.

However, it is crucial to know how different methods compare to each other, and to answer questions such as the following: how many relevant and irrelevant patterns does each algorithm find in relation to given reference data? Are there specific advantages of a string-based over a geometric approach for a given pattern discovery task, or vice versa? In short: which musical pattern discovery method performs best as compared against a given set of reference data?

3.2 Music representations

As of yet, there is no systematic analysis of the influence of the music representation on musical pattern discovery. The presented studies use various representations, but they do not directly evaluate the influence of the music representation on the musical pattern discovery results.

This is another challenge of musical pattern discovery: do more abstract representations lead to more irrelevant discovered patterns? Or is some abstraction desired for the description of some musical parameters? Which musical dimensions are important for a given musical genre? Which music representation approximates best the judgements on repeating patterns of a human listener, as manifested in reference data?

3.3 Filtering

We have also seen manifold filtering criteria applied in various studies. Again, the influence of different filtering parameters on the quality of musical pattern discovery has not yet been systematically investigated.

Are the most frequent patterns the ones human listeners would judge to be the most relevant? Or is it, conversely, just the unique patterns which capture our attention? Which filtering criteria match those of a human annotator of a given reference annotation most closely?

Insights on the influence of filtering are not only needed to improve musical pattern discovery methods; filtering criteria of a computational method can also serve as a model of the criteria human listeners apply when they isolate themes, motifs, or other salient musical patterns in musical pieces.

3.4 Reference data

As the preceding paragraphs show, there is a great need for references of human judgements on repeated patterns in music.

The literature overview pointed out a number of resources which can be used as reference data for musical pattern discovery. As this reference data is very diverse, and highly subjective, matching the results of a computational method closely to one set of reference data does not necessarily imply that the whole problem of musical pattern discovery has been solved.

This is another challenge of musical pattern discovery: will an algorithm which performs well on finding licks in jazz improvisation also perform well for finding themes in Western art music? What exactly *is* the data contained in a reference annotation?

3.5 Evaluation

Further advances in repeated melodic pattern discovery research hinges on a good evaluation of different methods. Many of the above studies report some successful findings, yet how do we know these findings are not only the grains the proverbial blind hen happens to find?

One very essential approach is of course a qualitative evaluation, in which discovered patterns are compared to those in reference data. However, for large corpora this is unfeasible, and it is hard to report the relative success of a method using a qualitative evaluation, unless the lengthy analysis would be reported.

We showed above that some studies use quantitative measures, derived from the number of agreements between algorithmically discovered patterns and patterns in reference data. Yet the definition of an agreement is problematic: do we only consider patterns which match those in a reference exactly, or should there also be a (penalized) score for partial matches?

Collins et al. suggest to allow a certain number of differences between algorithmically discovered and reference patterns [10]. A window around the start and end position of a pattern, allowing patterns to be counted as matches which are slightly shorter or longer than the reference patterns, might be another worthwhile approach. Recently, Collins [9] suggested a cardinality score as a similarity measure, which expresses the amount of the overlap between patterns. From this, several potentially interesting measures can be derived (see [9] for details).

So far, it is not clear whether some of these evaluation measures might be too tolerant, including too many patterns as agreements between algorithm and reference, or too strict, and register almost no patterns as matches between algorithm and reference. In both extremes, the performances of different methods, music representations and filtering techniques would probably be hard to assess.

Finding an effective, unified methodology of evaluation represents another challenge to the field of musical pattern discovery.

4 Perspectives of musical pattern discovery

From the above mentioned challenges, we conclude important perspectives for future research on musical pattern discovery. These perspectives relate to the influence different methods, different music representations and different filtering techniques have on the results of musical pattern discovery, what we can learn from annotated patterns in reference data, and which evaluation measures should be used.

4.1 Comparison of Methods

In order to assess different methods for repeated melodic pattern discovery, it is indispensable to compare them in a more direct way, even if they have been designed to discover specific kinds of patterns, or to discover patterns in a specific genre. We suggested in our overview a taxonomy of different approaches to pattern discovery, which helps to put the many different methods in perspective.

Ideally, next to the qualitative comparison of results, also a quantitative comparison should take place, which necessitates that methods be tested on the same corpus, and compared against the same reference data. The recently introduced MIREX track on musical pattern discovery [9] is an important step into this direction.

4.2 Comparison of music representations

We have noted the importance of understanding the influence of music representation on musical pattern discovery. Future research can amend this gap of knowledge by systematically comparing results from different representations to reference data. This way, we can learn more about the musical dimensions and abstraction levels that the human analysts who created reference data rely on for their pattern discovery.

Results from experimental studies (e.g. [15, 24, 42]) on music perception and recall should also be taken into account for choosing music representations. They provide theories which can be employed and tested by musical pattern discovery. The comparison of musical pattern discovery using different music representations with human annotations and with perceptual theories will generate insights which can feed back into research on similarity and variation in music theory and music cognition.

4.3 Comparison of filtering techniques

The systematic investigation of filtering techniques is also marked out as a fruitful direction of future research. By applying different filtering criteria and comparing the resulting patterns to reference data, we can better understand what conditions a pattern needs to fulfill in order to be considered relevant by human analysts.

Moreover, insights from research on long-term musical salience [2] can lead to models of how human listeners filter musical patterns according to salience or relevance. Musical pattern discovery can benefit from and contribute to this research area in music cognition.

4.4 Reference data

We have shown that there is a need for reference data of musical patterns annotated by human listeners. At the same time, the available reference data is a challenge in itself, as different references are based on subjective judgements of human analysts who annotated different kinds of patterns, working in different genres.

Researchers in musical pattern discovery can treat their algorithms as models of human analysis, and through relative successes and failures of these models, we can understand better which criteria underlie human judgements on repeated musical patterns.

Therefore, studying the reference data itself using musical pattern discovery is another perspective of this research area, which will promote knowledge on the concept of musical repetition, as applied by different human listeners in different genres. Such knowledge is important for various disciplines of music research interested in the nature of musical repetitions.

4.5 Towards an evaluation methodology

We infer that the best way to quantitatively evaluate musical pattern discovery algorithms consists of a combination of the several proposed measures. In combination, these different evaluation measures should give a reasonable impression of the respective successes of different approaches.

Many of the presented studies have performed a qualitative analysis of selected patterns. This should remain an indispensable evaluation step: patterns which quantitatively correspond to reference data exceptionally well, but also those which correspond exceptionally badly should be investigated qualitatively, to get a clearer view of the behaviour of a pattern discovery algorithm.

Quantitative and qualitative methods can also be combined in a learning process, by which quantitative measures are fitted to qualitative judgements. This method has been applied to improve the automatic generation of musical patterns by Pearce and Wiggins [38]. A use of qualitative judgements in combination with different quantitative measures will not only contribute to the improvement of musical pattern discovery, but will also greatly benefit related research in other fields, e.g. music cognition, musicology and music theory, as notions of what constitutes a repetition are quantified.

5 Conclusion

Our literature overview has shown that between the many approaches to musical pattern discovery, there are important challenges left unaddressed. There

is a need to investigate how methods developed for a specific goal generalise to other tasks within musical pattern discovery, and how different methods perform on the same reference data. Moreover, the influence of different music representations and filtering on the results of musical pattern discovery algorithms is not yet unravelled, which demands further research. The different idiosyncrasies of reference annotations, and their semantics are poorly understood. Finally, no standardized evaluation measure for musical pattern discovery has yet been established, which we suggest to overcome by combining various quantitative measures with a qualitative evaluation.

As stated in the introduction, much is to be gained for diverse music research disciplines from musical pattern discovery. Some successes have already been achieved, which makes further research in this field an intriguing effort to pursue. This thorough investigation of the state of the art in musical pattern discovery, its challenges and perspectives, should help to mark out the field in which further fruitful research can take place.

Acknowledgments. We thank Henkjan Honing and the reviewers for their useful comments, which greatly helped to improve this paper. Berit Janssen and Peter van Kranenburg are supported by the Computational Humanities Programme of the Royal Netherlands Academy of Arts and Sciences, under the auspices of the Tunes&Tales project. For further information, see <http://ehumanities.nl>. Anja Volk and W. Bas de Haas are supported by the Netherlands Organization for Scientific Research through an NWO-VIDI grant to Anja Volk (276-35-001).

References

1. Barlow, H., Morgenstern, S.: A dictionary of musical themes. Crown Publishers, New York (1948)
2. Burgoyne, J.A., Bountouridist, D., Van Balent, J., Honing, H.: Hooked: a Game for Discovering what Makes Music Catchy. In: Proceedings of the 14th International Society for Music Information Retrieval Conference (in press) (2013)
3. Buteau, C., Mazzola, G.: Motivic analysis according to Rudolph R eti: formalization by a topological model. *Journal of Mathematics and Music: Mathematical and Computational Approaches to Music Theory, Analysis, Composition and Performance* 2(3), 117–134 (2008)
4. Buteau, C., Viperman, J.: Melodic Clustering Within Motivic Spaces : Visualization in OpenMusic and Application to Schumanns Tr aumerei. In: Klouche, T., Noll, T. (eds.) *Mathematics and Computation in Music*, pp. 59–66. Springer, Berlin (2009)
5. Cambouropoulos, E.: Musical Parallelism and Melodic Segmentation: A Computational Approach. *Music Perception* 23(3), 249–268 (2006)
6. Cambouropoulos, E., Crochemore, M., Iliopoulos, C.S., Mohamed, M., Sagot, M.F.: All maximal-pairs in stepleap representation of melodic sequence. *Information Sciences* 177(9), 1954–1962 (May 2007), <http://linkinghub.elsevier.com/retrieve/pii/S0020025506003525>
7. Chou, T.c., Chen, A.L.P., Liu, C.c.: Music Databases : Indexing Techniques and Implementation. In: *Multimedia Database Management Systems*. pp. 46–53 (1996)

8. Clifford, R., Iliopoulos, C.: Approximate string matching for music analysis. *Soft Computing* 8, 597–603 (2004)
9. Collins, T.: Discovery of repeated themes and sections. http://www.music-ir.org/mirex/wiki/2013:Discovery_of_Repeated_Themes_%26_Sections, section 4, accessed: 2013-05-04
10. Collins, T., Thurlow, J., Laney, R., Willis, A., Garthwaite, P.H.: A comparative evaluation of algorithms for discovering translational patterns in Baroque keyboard works Conference Item. In: *International Symposium for Music Information Retrieval* (2010)
11. Conklin, D.: Multiple Viewpoint Systems for Music Prediction. *Journal of New Music Research* 24(1), 51–73 (1995)
12. Conklin, D.: Antipattern Discovery in Folk Tunes. *Journal of New Music Research* 42(2), 161–169 (Jun 2013), <http://www.tandfonline.com/doi/abs/10.1080/09298215.2013.809125>
13. Conklin, D., Anagnostopoulou, C.: Comparative Pattern Analysis of Cretan Folk Songs. *Journal of New Music Research* 40(2), 119–125 (Jun 2011), <http://www.tandfonline.com/doi/abs/10.1080/09298215.2011.573562>
14. Crochemore, M.: An Optimal Algorithm for Computing the Repetitions in a Word. *Information Processing Letters* 12(5), 244–250 (1981)
15. Dowling, W.J.: Scale and Contour: Two Components of a Theory of Memory for Melodies. *Psychological Review* 85(4), 341–354 (1978)
16. Giraud, M., Groult, R., Levé, F.: Subject and counter-subject detection for analysis of the Well-Tempered Clavier fugues. In: *International Symposium on Computer Music Modelling and Retrieval (CMMR)*. vol. 2012 (2012)
17. Gusfield, D.: *Algorithms on strings, trees and sequences: computer science and computational biology*. Cambridge University Press (1997)
18. Honing, H.: Computational Modeling of Music Cognition: A Case Study on Model Selection. *Music Perception* 23(5), 365–376 (2006)
19. Honing, H.: Lure(d) into listening : The potential of cognition-based music information retrieval. *Empirical Musicology Review* 5(4), 146–151 (2010)
20. Huron, D.: *Sweet Anticipation. Music and the Psychology of Expectation*. MIT Press, Cambridge, Massachusetts (2007)
21. Janssen, B., Honing, H., Kranenburg, P.v., Grijp, L.: Stability of melodic patterns in oral transmission. In: *Proceedings of the Society for Music Perception and Cognition* (2013)
22. Karydis, I., Nanopoulos, A., Manolopoulos, Y.: Finding maximum-length repeating patterns in music databases. *Multimedia Tools and Applications* 32, 49–71 (Oct 2007), <http://link.springer.com/10.1007/s11042-006-0068-5>
23. Klapuri, A.: Pattern induction and matching in music signals. In: *Exploring Music Contents. 7th International Symposium, CMMR 2010, Málaga, Spain*. pp. 188–204 (2010)
24. Klusen, E., Moog, H., Piel, W.: Experimente zur mündlichen Tradition von Melodien. In: *Jahrbuch für Volksliedforschung*, pp. 11–23 (1978)
25. Knopke, I., Jürgensen, F.: A System for Identifying Common Melodic Phrases in the Masses of Palestrina. *Journal of New Music Research* 38(2), 171–181 (Jun 2009), <http://www.tandfonline.com/doi/abs/10.1080/09298210903288329>
26. Knuth, D.E., Morris, J.H., Pratt, V.R.: Fast Pattern Matching in Strings. *SIAM Journal of Computing* 6(2), 323–350 (1977)
27. Koniari, D., Predazzer, S., Mélen, M.: Categorization and Schematization Processes Used in Music Perception by 10- to 11-Year-Old Children. *Music Perception: An Interdisciplinary Journal* 18(3), 297–324 (2001)

28. Lartillot, O.: Multi-dimensional motivic pattern extraction founded on adaptive redundancy filtering. *Journal of New Music Research* 34(4), 375–393 (Dec 2005), <http://www.tandfonline.com/doi/abs/10.1080/09298210600578246>
29. Lee, W., Chen, A.L.P.: Efficient Multi-Feature Index Structures for Music Data Retrieval. In: *Electronic Imaging*. International Society for Optics and Photonics. pp. 177–188 (1999)
30. Lemström, K., Laine, P.: Music Information Retrieval Using Musical Parameters. In: *International Computer Music Conference*. vol. 26, pp. 341–348 (1998)
31. Levenshtein, V.I.: Binary Codecs Capable of Correcting Deletions, Insertions and Reversals. *Soviet Physics - Doklady* 10(8), 707–710 (1966)
32. Marsden, A.: Counselling a better relationship between mathematics and musicology. *Journal of Mathematics and Music: Mathematical and Computational Approaches to Music Theory, Analysis, Composition and Performance* 6(2), 145–153 (2012)
33. Meek, C., Birmingham, W.P.: Thematic Extractor. In: *International Symposium for Music Information Retrieval*. pp. 119–128 (2001)
34. Meredith, D., Lemström, K., Wiggins, G.A.: Algorithms for discovering repeated patterns in multidimensional representations of polyphonic music. *Journal of New Music Research* 31(4), 321–345 (2002)
35. Nieto, O., Farbood, M.M.: Perceptual Evaluation of Automatically Extracted Musical Motives. In: *Proceedings of the 12th International Conference on Music Perception and Cognition*. pp. 723–727 (2012)
36. Owens, T.: *Charlie Parker : techniques of improvisation*. Ph.D. thesis (1974)
37. Paulus, J., Müller, M., Klapuri, A.: State of the art report: Audio-based music structure analysis. In: *Proc. ISMIR*. pp. 625–36 (2010)
38. Pearce, M.T., Wiggins, G.A.: Evaluating Cognitive Models of Musical Composition. In: *Proceedings of the 4th international joint workshop on computational creativity*. pp. 73–80. London (2007)
39. Repp, B.H.: Diversity and Commonality in Music Performance: An Analysis of Timing Microstructure in Schumann' sTräumerei. *The Journal of the Acoustical Society of America* 92, 227–260 (1992)
40. Réti, R.: *The thematic process in music*. Macmillan, New York (1951)
41. Rolland, P.Y.: Discovering Patterns in Musical Sequences. *Journal of New Music Research* 28(4), 334–351 (1999)
42. Sloboda, J., Parker, D.: Immediate recall of melodies. In: Howell, P., Cross, I., West, R. (eds.) *Musical structure and cognition*, pp. 143–167. Academic Press, London (1985)
43. Szeto, W.M., Wong, M.H.: A graph-theoretical approach for pattern matching in post-tonal music analysis. *Journal of New Music Research* 35(4), 307–321 (Dec 2006), <http://www.tandfonline.com/doi/abs/10.1080/09298210701535749>
44. Volk, A., de Haas, W.B., van Kranenburg, P.: Towards modelling variation in music as foundation for similarity. In: *Proc. ICMPC*. pp. 1085–1094 (2012)
45. Volk, A., van Kranenburg, P.: Melodic similarity among folk songs: An annotation study on similarity-based categorization in music. *Musicae Scientiae* 16(3), 317–339 (Jul 2012), <http://msx.sagepub.com/lookup/doi/10.1177/1029864912448329>
46. Wittlich, G.: Interval Set Structure in Schoenberg's op. 11, no. 1. *Perspectives of New Music* pp. 41–55 (1974)