

Finding 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

{berit.janssen,peter.van.kranenburg}@meertens.knaw.nl

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

³ Department of Information and Computing Sciences,
Utrecht University, Utrecht, The Netherlands
{w.b.dehaas,a.volk}@uu.nl

Abstract. This paper discusses the current state of knowledge on musical pattern finding. Various studies propose computational methods to find repeated musical patterns. Our detailed review of these studies reveals important challenges in musical pattern finding 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 finding: 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 identified 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 finding · Musical pattern discovery · Musical pattern matching · 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 [25, 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 finding is important in several areas of music research. In Music Information Retrieval, repetitions have been

used as indicators of musical segmentation [21], or to find themes or choruses in large databases [49]. In Music Analysis, analytical approaches based on repetition, for instance Réti's motivic analysis [52], have been formalized and evaluated by developing a computer model [5]. 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 [58]. Hypotheses on memory, recall and transmission of melodies can be tested on large databases using musical pattern finding [26].

As we show in this article, there are many different kinds of repetition that different researchers investigate: large repeated structures, such as themes, choruses, 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 formalized 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 [57].

A computational method to find 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 [22,23]. In this paper, we explore computational methods for finding repeated musical patterns, which we will refer to as musical pattern finding.

Currently, the knowledge of musical pattern finding 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 finding is currently unknown. This lack of comparative assessment is further complicated by the lack of a sound methodology for the evaluation of musical pattern finding.

This paper provides a comprehensive overview, review and discussion of the field of musical pattern finding. We present the essence of assorted studies, and we proceed to clarify the relationships between different methods, proposing a taxonomy of musical pattern finding 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 finding and conclude with steps to overcome these challenges.

2 State of Knowledge on Musical Pattern Finding

We conducted a comprehensive literature survey on musical pattern finding. 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.

In this survey, our focus lies on studies using symbolic music representations. Several studies work with audio representations [9,44], but we will not discuss the problems related to transforming audio representations for musical pattern finding. For this, and other methods for the audio domain, we refer the reader to the overview by Klapuri [28].

We present our overview of studies on musical pattern finding in Table 1. The tables' 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 finding. 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 Table 1, the interest for musical pattern finding 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 most of the other research on musical pattern finding.

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 Fig. 1. Below, we will explain the distinctions within the taxonomy.

Pattern Discovery or Pattern Matching. Musical pattern discovery, on the one hand, aims at the identification of motifs, themes, and other musical structures in a piece of music, or between related pieces of music (intra-, and inter-opus discovery). Typically, algorithms applied for pattern discovery do not presuppose prior knowledge of possible candidates.

Musical pattern matching, on the other hand, strives to recognize pre-defined patterns in a musical piece, or a corpus of music [53]. Applications of pattern matching include the identification of fugue subjects and counter-subjects in Bach fugues [19] or variations of themes, as in some works by Mozart [18], among others.

In terms of the applied techniques, there is an overlap of pattern matching with classification problems in music retrieval, which aim at the identification of similar musical pieces in a large database (e.g. [24,33]). However, in this paper,

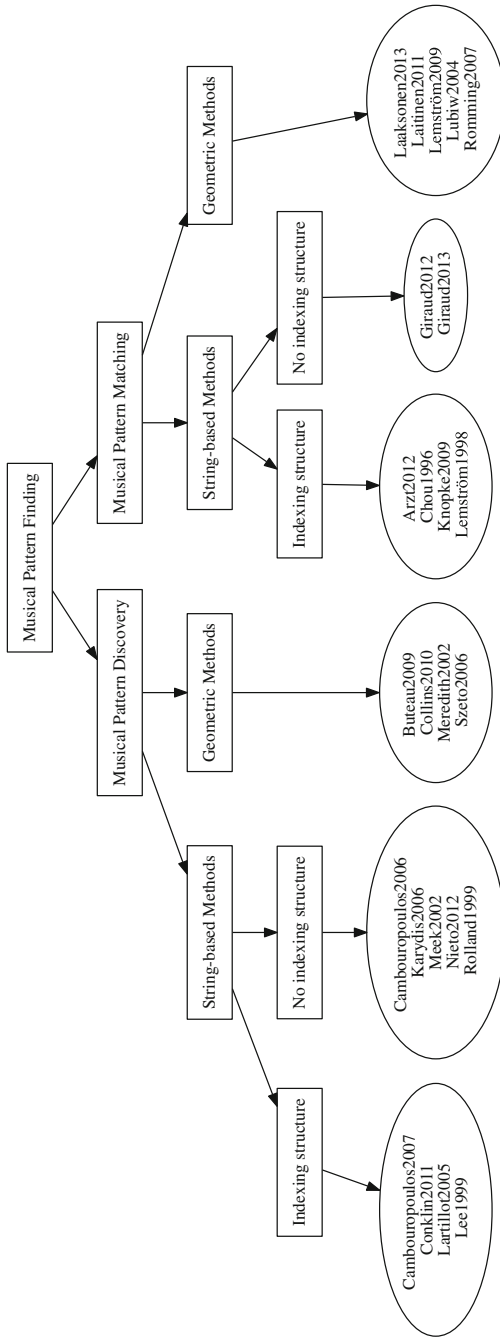


Fig. 1. A schematic representation of a taxonomy for musical pattern finding methods. For each category, relevant studies are listed.

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

Study	Goals	Musical pieces	Method	Music representation	Filtering	Reference data	Evaluation
Arzt2012 [1]	Fast identification of piece and score position from a performance	Selected pieces by Mozart and Chopin	Generation of fingerprints by sampling pitches at fixed time intervals, which are matched as hash keys	pitch, inter onset intervals		Performance data of musical pieces	Precision of results in terms of piece and score position
Bateau2009 [6]	Computer-aided motivic analysis	Schumann's <i>Träumerei</i>	topological model: motifs as shapes in n-dimensional space	pitch and onset	shape clustering	Repp1992 [51]	Qualitative comparison with music theorist. Repp's segmentation of Schumann's <i>Träumerei</i>
Cambouropoulos2006 [7]	Segmentation based on musical repetitions	Sonatina No. 2 in C Major by Anton Diabelli (soprano voice)	set partitioning method[16]	pitch intervals; step-leap	frequency of occurrence, pattern length	Koniari2001 [32]	Quantitative comparison with human segmentations of a sonatina by Anton Diabelli
Cambouropoulos2007 [8]	Finding motifs in musical pieces		Searching maximal pairs in suffix trees	pitch interval, step-leap representation of pitch			
Chou1996 [9]	Content-based retrieval of music from large data bases		Suffix trees	representation of chords			
Collins2010 [12]	Finding translational patterns in musical pieces to support music analysis	two sonatas by Scarlatti, two Bach preludes (polyphonic)	Structure Induction Algorithms (extension of Meredith2002)	pitch interval, onset	<i>compactness</i> patterns with adjacent data-points	Collins2010 [12]	Quantitative comparison with motif annotations in two sonatas by Scarlatti and two Bach preludes; using Precision and Recall
Conklin2011 [15]	Classifying Cretan folk songs based on characteristic patterns	106 monophonic Cretan folksongs	sequential pattern mining: finding maximally general distinctive patterns (MGDP) in a search tree	pitch intervals	frequency of occurrence		Qualitative analysis of patterns as parts of Cretan folk music genres

(Continued)

Table 1. (*Continued*)

Study	Goals	Musical pieces	Method	Music representation	Filtering	Reference data	Evaluation
Giraud2012 [19]	Finding fugue subjects	24 fugues from Bach's Wohltemperiertes Clavier	Alignment of a fugue subject with the remaining fugue	(interval, onset, length), (pitch, onset, length)	Threshold of similarity	Bruhn1993 [3]	Qualitative; precision and recall
Giraud2013 [18]	Finding variations of a theme	Two works by W.A. Mozart: K. 331 and K. 265	Alignment of theme and variation using fragmentation	pitch, onset, length	length, at least one pitch equal, equivalent position	Musical form	Precision, recall
Karydis2006 [27]	Finding themes in large music databases	Assorted MIDI files (unspecified)	Finding maximum-length repeating patterns with set partitioning method	pitch			Computation speed
Knopke2009 [30]	Analysis of Palestrina's contrapuntal structure	masses by Palestrina (separated into voices and phrases)	4 suffix arrays for each phrase (indicated by rests), for original, inversion, retrograde, retrograde inversion	pitch and duration intervals, mapped to UTF-8 tokens			Qualitative analysis of selected patterns, especially longest patterns and most frequent patterns
Laaksonen2013 [34]	Finding time-scaled and time-warped permutations of fugue subjects	unspecified music databases	Dynamic time warping	onset times, pitches			Performance speed
Laitinen2011 [35]	Finding more efficient algorithms for exact and approximate matches, time-scaled and time-warped	Music data from Musitopia database (unspecified)	Dynamic time warping	pitch against time			Performance speed

Table 1. (Continued)

Study	Goals	Musical pieces	Method	Music representation	Filtering	Reference data	Evaluation
Lartillot2005 [36]	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 patterns, filtered during the tree construction		Qualitative analysis of selected patterns
Lee1999 [37]	Content-based retrieval of music from large data bases	100 songs in MIDI format (unspecified)	Different suffix tree algorithms	pitch interval, quantized duration (combined and separate)			Computation speed
Lenström1998 [38]	Translation-invariant retrieval from music databases	154 pieces of music from different genres	Suffix tree	pitch/duration tuples for different abstraction levels			occurrences of search keys, and quantity of melodies containing search keys
Lenström2009 [39]	Content-based music retrieval of motifs and themes	Polyphonic music (not specified)	defining vectors connecting lexicographically ordered points within patterns, detecting them in a large database	pitch against time	comparison of 7 different filters for exact and approximate matching		Precision and recall of patterns artificially inserted into the database
Lubiw2004 [42]	Finding matches of musical patterns translated in pitch and time	Bach invention, Bach fugue, Mozart sonata	Alignment with weights following Mongeau and Sankoff [46]	onset, offset, pitch			Qualitative analysis of patterns
Meek2001 [44]	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 [2]	Quantitative comparison of selected examples against <i>Dictionary of musical themes</i> (Barlow): the evaluation procedure is not described

(Continued)

Table 1. (*Continued*)

Study	Goals	Musical pieces	Method	Music representation	Filtering	Reference data	Evaluation
Meredith2002 [45]	Assisting music analysts and composers by finding repeated patterns in music	52 polyphonic pieces, mostly from Bach's <i>Wohlttem-periertes 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 [47]	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 [53]	Automatizing pattern 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 [48]	Qualitative comparison of selected patterns with Owens' catalogue of Charlie Parker's patterns
Romming2007 [54]	Finding matches of musical patterns in polyphonic music, shifted in pitch and time and meter-invariant	"Fuggi, fuggi, da questo cielo" with variations; Bach's Goldberg Variation No.14	Comparison of alignment based on Hausdorff metric and geometric hashing	pitch, duration, onset			Qualitative analysis
Szeto2006 [56]	Discovering structures in post-tonal music	Arnold Schoenberg's <i>Klavierstücke</i> , op. 11 no.1	Graph with notes as vertices, edges linking simultaneous and sequential events	pitch class, onset	Clustering of graph vertices	Wittlich1974 [59]	Quantitative comparison with musicological analysis (Wittlich1974): counting number of matches and approximate matches

we focus on the problem of identifying a musical subsequence in a larger piece of music.

In Fig. 1, musical pattern discovery algorithms are presented on the left; musical pattern matching approaches can be found on the right side of the diagram.

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 finding 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 [20] provides a thorough overview of these techniques.

In geometric methods, a melody is considered as a shape in an n -dimensional space. Repeated patterns are then identified as (near-)identical shapes. Geometric methods are especially interesting for polyphonic music, as they can deal more efficiently with note events occurring at the same time [45, p. 328]. Meredith's *Structure Induction Algorithms* (SIA) [45] order the pitch/onset points lexicographically, and search for vectors of pitch and time relationships which repeat elsewhere in a musical piece. This concept, used by Meredith and Collins [12, 45] for pattern discovery, has been applied by Lemström et al. [39] for pattern matching.

The musical pattern discovery approaches by Buteau and Viperman [6] and Szeto and Wong [56] rely on a similar conceptualisation of music as n -dimensional shapes. In the latter study, the geometric relationships are represented as nodes and edges in graphs [56]. Laaksonen [34], Laitinen and Lemström [35], Lubiw and Tanur [42] and Romming and Selfridge-Field [54] represent musical pieces as point sets of onset times and pitch, which can be shifted in the time or pitch domain, which is handled through substitution weights in an alignment algorithm.

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.

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 [47]. 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 [31], has been applied by Rolland [53] for musical pattern discovery.

Yet another approach is Crochemore's set partitioning method [16], which recursively splits the melody M into sets of repeating tokens. Cambouropoulos [7] 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. [27] 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 [44] 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 [8, 9, 15, 30, 36, 37]. Knopke and Jürgensen [30] use suffix arrays representing phrases of Palestrina masses. Conklin and Anagnostopoulou [15] use a sequence tree to represent search spaces of patterns in Cretan folk songs, which is pruned based on the patterns' frequency of occurrence. Lemström and Laine compare suffix tries to the more compressed form of suffix trees on a varied music corpus [38]. Lartillot [36] 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 finding. 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.

Several ways to define approximate matching for musical pattern finding have been proposed. Approximate matching algorithms typically distinguish between approximate matches and irrelevant patterns based on a threshold on a given similarity measure. This can be the number of allowed mismatches, referred to as *k-mismatch*, the Hausdorff-distance, which refers to the Euclidian distance between sets of points [54], or the length of the longest common subsequence [40]. If musical pieces are represented as sequences of numbers (e.g. pitch values), it is also possible to define a threshold on the difference between values of two compared sequences – δ -matching – or on the sum of all differences between the values in two compared strings – γ -matching – as suggested by Clifford and Iliopoulos [10]. However, to our knowledge, δ -matching and γ -matching have not been used for musical pattern finding to this date.

Furthermore, the threshold can also be defined as a maximum amount of edit operations in an alignment algorithm (also known as edit distance or Levenshtein

distance) [41]. This way, also strings of different length, or strings containing gaps in relation to each other, can be considered as approximate matches. Giraud et al. [19] use edit distance as a similarity measure for Bach fugues. For the comparison of Mozart themes and variations, Giraud and Cambouropoulos [18] extend the alignment metric with another operation, fragmentation. Fragmentation was introduced by Mongeau and Sankoff [46] so that musical embellishments cause smaller edit distances between patterns than with the classical definition of edit distance.

In a generalisation for patterns repeated on different time scales, alignment algorithms have also been developed further so as to consider patterns shifted by time constants or ratios. This approach, generally referred to as dynamic time warping, has been applied, among others, by Laitinen and Lemström [35] to a database of music from different genres, and by Laaksonen [34] in order to find augmentations and diminutions of fugue subjects.

To our knowledge, only Rolland [53] applied approximate matching to musical pattern discovery, using the Levenshtein distance to compare a pattern with a match candidate. Implicitly, however, approximate matching can also be achieved through more abstract music representations, as Cambouropoulos et al. [8] 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. Moreover, 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. This is closely related to Conklin's [13] notion of musical viewpoints.

A glance at the music representation column reveals that the majority of the studies on musical pattern finding 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. [9] extract 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. Similarly, Arzt et al. [1] match patterns with a symbolic fingerprint that is derived from pitch values sampled at specific time intervals.

Three studies [8, 37, 53] suggest multiple music representations. Rolland [53] 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. [8] 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 [37] 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 patterns 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.” [43].

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 [7] filters according to pattern lengths. Likewise, Collins’ *compactness trawling* to refine the result of Structure Induction Algorithms [12] is based on such a filtering, in which patterns of adjacent notes are preferred over patterns with many intervening notes.

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 [7, 53].

Conklin and Anagnostopoulou [15] 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 analyzed piece. Conversely, also patterns which are characteristically not represented in specific pieces or genres can be interesting for music researchers [14].

Nieto and Farbod filter according to Gestalt rules during the search process, which means that pattern candidates containing relatively long notes or rests, or relatively large intervals will be rejected [47].

Lartillot’s detection of maximally specific patterns [36] filters for patterns which occur 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 optimisation algorithm. Meek and Birmingham [44] 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* [2].

Lemström et al. [39] apply seven different filtering methods to their pattern matching results: these filters are based on the relationships between the points in the query patterns (*intra-pattern vectors*). The least frequent intra-pattern vectors are considered the best candidates for meaningful matches.

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 [48] to themes in Western art music [2]. This reference data is typically assembled by domain specialists, annotating what they consider the most relevant patterns of the analyzed music collection.

Such reference data can serve as a ground truth 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

Many studies discussed here evaluate qualitatively: the discovered patterns are scrutinized. 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 [12, 18, 19, 44, 47, 56].

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 [6, 7].

In a number of the studies [34, 35, 37], the researchers aimed for fast solutions, which make the algorithms more interesting for practical use. Therefore, computation speed is used as an evaluation metric in these cases. This does not give an indication of the usefulness of the automatically found patterns, however.

Several quantitative evaluation measures have been suggested for musical pattern finding, but there seems to be no common methodology among 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 Finding

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 finding 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 finding task, or vice versa? In short: which musical pattern finding method performs best as compared against a given set of reference data?

3.2 Music Representation

As of yet, there is no systematic analysis of the influence of the music representation on musical pattern finding. 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 finding: 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 finding 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 finding 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 finding. 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 finding has been solved.

This is another challenge of musical pattern finding: will an algorithm that 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 finding research hinges on a good evaluation of different methods. Many of the above studies report some successful results, yet how do we know these results 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 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 found 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 or matched, and reference patterns [12]. 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 [11] 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 [11] 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 finding.

4 Perspectives of Musical Pattern Finding

From the above mentioned challenges, we conclude important perspectives for future research on musical pattern finding. 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 finding, it is indispensable to compare them in a more direct way, even if they have been designed to find specific kinds of patterns, or to find patterns in a specific genre. We suggested in our overview a taxonomy of different approaches to pattern finding, 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 [11] is an important step into this direction. Similar incentives to compare approaches in musical pattern matching would be highly desirable.

4.2 Comparison of Music Representations

We have noted the importance of understanding the influence of music representation on musical pattern finding. 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, abstraction levels and concepts of similarity that human analysts rely on for finding musical patterns.

Results from experimental studies (e.g. [17, 29, 55]) 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 finding. The comparison of musical pattern finding 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 [4] can lead to models of how human listeners filter musical patterns according to salience or relevance. Musical pattern finding 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 finding 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, another perspective of this research area is the study of reference data itself by comparison to computational results. Such a study 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 finding 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: for musical pattern discovery, patterns which quantitatively correspond to reference data exceptionally well, but also those which correspond exceptionally badly should be investigated qualitatively. In musical pattern matching, those patterns should be investigated which are found by the algorithm, but have not been annotated by human listeners, and those which have been annotated, but were overseen by the algorithm. Only with such an accompanying qualitative analysis can the behaviour of a pattern finding algorithm be truly understood.

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 [50]. A use of qualitative judgements in combination with different quantitative measures will not only contribute to the improvement of musical pattern finding, 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 finding, there are important challenges left unaddressed. There is a need to investigate how methods developed for a specific goal generalize to other tasks within musical pattern finding, 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 finding 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 finding 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 finding. 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 finding, 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 Organisation for Scientific Research through an NWO-VIDI grant to Anja Volk (276-35-001).

References

1. Arzt, A., Sebastian, B., Widmer, G.: Fast identification of piece and score position via symbolic fingerprinting. In: International Symposium for Music Information Retrieval (2012)
2. Barlow, H., Morgenstern, S.: A Dictionary of Musical Themes. Crown Publishers, New York (1948)
3. Bruhn, S.: J. S. Bach's Well-Tempered Clavier: In-depth Analysis and Interpretation. Mainer, Hong Kong (1993). <http://www-personal.umich.edu/siglind/text.htm>
4. Burgoyne, J.A., Bountouridis, D., Van Balen, J., Honing, H.: Hooked: a game for discovering what makes music catchy. In: Proceedings of the 14th International Society for Music Information Retrieval Conference (2013)
5. Buteau, C., Mazzola, G.: Motivic analysis according to Rudolph R eti: formalization by a topological model. *J. Math. Music Math. Comput. Approaches Music Theory Anal. Compos. Perform.* **2**(3), 117–134 (2008)
6. Buteau, C., Viperman, J.: Melodic clustering within motivic spaces : visualization in OpenMusic and application to Schumann's Tr aumerei. In: Klouche, T., Noll, T. (eds.) *Mathematics and Computation in Music*, pp. 59–66. Springer, Berlin (2009)
7. Cambouropoulos, E.: Musical parallelism and melodic segmentation: a computational approach. *Music Percept.* **23**(3), 249–268 (2006)
8. Cambouropoulos, E., Crochemore, M., Iliopoulos, C.S., Mohamed, M., Sagot, M.F.: All maximal-pairs in step-leap representation of melodic sequence. *Inf. Sci.* **177**(9), 1954–1962 (2007). <http://linkinghub.elsevier.com/retrieve/pii/S0020025506003525>
9. 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)
10. Clifford, R., Iliopoulos, C.: Approximate string matching for music analysis. *Soft Comput.* **8**, 597–603 (2004)

11. Collins, T.: Discovery of repeated themes and sections, Sect. 4. http://www.music-ir.org/mirex/wiki/2013:Discovery_of_Repeated_Themes_%26_Sections. Accessed 4 May 2013
12. 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)
13. Conklin, D.: Multiple viewpoint systems for music prediction. *J. New Music Res.* **24**(1), 51–73 (1995)
14. Conklin, D.: Antipattern discovery in folk tunes. *J. New Music Res.* **42**(2), 161–169 (2013). <http://www.tandfonline.com/doi/abs/10.1080/09298215.2013.809125>
15. Conklin, D., Anagnostopoulou, C.: Comparative pattern analysis of cretan folk songs. *J. New Music Res.* **40**(2), 119–125 (2011). <http://www.tandfonline.com/doi/abs/10.1080/09298215.2011.573562>
16. Crochemore, M.: An optimal algorithm for computing the repetitions in a word. *Inf. Process. Lett.* **12**(5), 244–250 (1981)
17. Dowling, W.J.: Scale and contour: two components of a theory of memory for melodies. *Psychol. Rev.* **85**(4), 341–354 (1978)
18. Giraud, M., Cambouropoulos, E.: Fragmentations with pitch, rhythm and parallelism constraints for variation matching. In: Proceedings of the 10th International Symposium on Computer Music Multidisciplinary Research, pp. 241–252 (2013)
19. 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)
20. Gusfield, D.: Algorithms on Strings, Trees and Sequences: Computer Science and Computational Biology. Cambridge University Press, Cambridge (1997)
21. de Haas, W.B., Volk, A., Wiering, F.: Structural segmentation of music based on repeated harmonies. In: Proceedings of IEEE International Symposium on Multimedia (ISM2013), pp. 255–258 (2013)
22. Honing, H.: Computational modeling of music cognition: a case study on model selection. *Music Percept.* **23**(5), 365–376 (2006)
23. Honing, H.: Lure(d) into listening : the potential of cognition-based music information retrieval. *Empir. Musicology Rev.* **5**(4), 146–151 (2010)
24. Huang, T., Xia, G., Ma, Y., Dannenberg, R., Faloutsos, C.: MidiFind: fast and effective similarity searching in large MIDI databases. In: Proceedings of the 10th International Symposium on Computer Music Multidisciplinary Research, pp. 209–224 (2013)
25. Huron, D.: Sweet Anticipation: Music and the Psychology of Expectation. MIT Press, Cambridge (2007)
26. Janssen, B., Honing, H., van Kranenburg, P., Grijp, L.: Stability of melodic patterns in oral transmission. In: Proceedings of the Society for Music Perception and Cognition (2013)
27. Karydis, I., Nanopoulos, A., Manolopoulos, Y.: Finding maximum-length repeating patterns in music databases. *Multimed. Tools Appl.* **32**, 49–71 (2007). <http://link.springer.com/10.1007/s11042-006-0068-5>
28. Klapuri, A.: Pattern induction and matching in music signals. In: Exploring Music Contents, 7th International Symposium, CMMR 2010, Málaga, Spain, pp. 188–204 (2010)
29. Klusen, E., Moog, H., Piel, W.: Experimente zur mündlichen Tradition von Melodien. In: Jahrbuch für Volksliedforschung, pp. 11–23 (1978)

30. Knopke, I., Jürgensen, F.: A system for identifying common melodic phrases in the masses of Palestrina. *J. New Music Res.* **38**(2), 171–181 (2009). <http://www.tandfonline.com/doi/abs/10.1080/09298210903288329>
31. Knuth, D.E., Morris, J.H., Pratt, V.R.: Fast pattern matching in strings. *SIAM J. Comput.* **6**(2), 323–350 (1977)
32. Koniari, D., Predazzer, S., Mélen, M.: Categorization and schematization processes used in music perception by 10- to 11-year-old children. *Music Percept. Interdisc. J.* **18**(3), 297–324 (2001)
33. Kranenburg, P.V.: A computational approach to content-based retrieval of Folk song melodies. Ph.D. thesis, University of Utrecht (2010)
34. Laaksonen, A.: Efficient and simple algorithms for time-scaled and time-warped music search. In: *Proceedings of the 10th International Symposium on Computer Music Multidisciplinary Research*, pp. 621–630, No. 118653 (2013)
35. Laitinen, M., Lemström, K.: Dynamic programming in transposition and time-warp invariant polyphonic content-based music retrieval. In: *12th International Society for Music Information Retrieval Conference (ISMIR 2011)*, pp. 369–374. *Ismir* (2011)
36. Lartillot, O.: Multi-dimensional motivic pattern extraction founded on adaptive redundancy filtering. *Journal of New Music Research* **34**(4), 375–393 (2005). <http://www.tandfonline.com/doi/abs/10.1080/09298210600578246>
37. Lee, W., Chen, A.L.P.: Efficient multi-feature index structures for music data retrieval. In: *Electronic Imaging*, pp. 177–188. *International Society for Optics and Photonics* (1999)
38. Lemström, K., Laine, P.: Music information retrieval using musical parameters. In: *International Computer Music Conference*, vol. 26, pp. 341–348 (1998)
39. Lemström, K., Mikkilä, N., Mäkinen, V.: Filtering methods for content-based retrieval on indexed symbolic music databases. *Inf. Retr.* **13**(1), 1–21 (2009). <http://link.springer.com/10.1007/s10791-009-9097-9>
40. Lemström, K., Ukkonen, E.: Including interval encoding into edit distance based music comparison and retrieval. In: *Proceedings of the AISB'2000 Symposium on Creative and Cultural Aspects and Applications of AI and Cognitive Science*, Birmingham (2000)
41. Levenshtein, V.I.: Binary codecs capable of correcting deletions, insertions and reversals. *Sov. Phys. - Dokl.* **10**(8), 707–710 (1966)
42. Lubiw, A., Tanur, L.: Pattern matching in polyphonic music as a weighted geometric translation problem. In: *International Society for Music Information Retrieval Conference (ISMIR 2004)*, pp. 369–374 (2004)
43. Marsden, A.: Counselling a better relationship between mathematics and musicology. *J. Math. Music Math. Comput. Approaches Music Theory Anal. Compos. Perform.* **6**(2), 145–153 (2012)
44. Meek, C., Birmingham, W.P.: Thematic extractor. In: *International Symposium for Music Information Retrieval*, pp. 119–128 (2001)
45. Meredith, D., Lemström, K., Wiggins, G.A.: Algorithms for discovering repeated patterns in multidimensional representations of polyphonic music. *J. New Music Res.* **31**(4), 321–345 (2002)
46. Mongeau, M., Sankoff, D.: Comparison of musical sequences. *Comput. Humanit.* **24**, 161–175 (1990)
47. 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)

48. Owens, T.: Charlie Parker: techniques of improvisation. Ph.D. thesis (1974)
49. Paulus, J., Müller, M., Klapuri, A.: State of the art report: audio-based music structure analysis. In: Proceedings of ISMIR, pp. 625–636 (2010)
50. Pearce, M.T., Wiggins, G.A.: Evaluating cognitive models of musical composition. In: Proceedings of the 4th International Joint Workshop on Computational Creativity, London, pp. 73–80 (2007)
51. Repp, B.H.: Diversity and commonality in music performance: an analysis of timing microstructure in schumann's träumerei. *J. Acoust. Soc. Am.* **92**, 227–260 (1992)
52. Réti, R.: *The Thematic Process in Music*. Macmillan, New York (1951)
53. Rolland, P.Y.: Discovering patterns in musical sequences. *J. New Music Res.* **28**(4), 334–351 (1999)
54. Romming, C.A., Selfridge-Field, E.: Algorithms for polyphonic music retrieval: the hausdorff metric and geometric hashing. In: International Society for Music Information Retrieval Conference (ISMIR 2007), pp. 457–462 (2007)
55. 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)
56. Szeto, W.M., Wong, M.H.: A graph-theoretical approach for pattern matching in post-tonal music analysis. *J. New Music Res.* **35**(4), 307–321 (2006). <http://www.tandfonline.com/doi/abs/10.1080/09298210701535749>
57. Volk, A., de Haas, W.B., van Kranenburg, P.: Towards modelling variation in music as foundation for similarity. In: Proceedings of ICMPC, pp. 1085–1094 (2012)
58. Volk, A., van Kranenburg, P.: Melodic similarity among folk songs: an annotation study on similarity-based categorization in music. *Musicae Sci.* **16**(3), 317–339 (2012). <http://msx.sagepub.com/lookup/doi/10.1177/1029864912448329>
59. Wittlich, G.: Interval set structure in schoenberg's op. 11, no. 1. *Perspect. New Music* **13**, 41–55 (1974)