

Computational Fugue Analysis

Abstract: One of the pinnacles of form in classical Western music, the fugue is often used in the teaching of music analysis and composition. Fugues alternate between instances of a subject and other patterns and modulatory sections, called episodes. Musicological analyses are generally built on these patterns and sections.

We have developed several algorithms to perform an automated analysis of a fugue, starting from a score in which all the voices are separated. By focusing on the diatonic similarities between pitch intervals, we detect subjects and countersubjects, as well as partial harmonic sequences inside the episodes. We also implemented tools to detect subject scale degrees, cadences, and pedals, as well as a method for segmenting the fugue into exposition and episodic parts.

Our algorithms were tested on a corpus of 36 fugues by J. S. Bach and Dmitri Shostakovich. We provide formalized ground-truth data on this corpus as well as a dynamic visualization of the ground truth and of our computed results. The complete system showed acceptable or good results for about one half of the fugues tested, enabling us to depict their design.

A fugue is a polyphonic musical piece built on several melodic themes that consist of a *subject* and, in most cases, one or more *countersubjects*. The Italian word *fuga* is related to the Latin words *fugere* (to flee) and *fugare* (to chase): The patterns are played by each voice, one following the other in succession. They occur either in their initial form or, more often, altered or transposed, building a complex harmonic texture. Many composers wrote fugues or included fugal parts in larger pieces. The two books of J. S. Bach's *Well-Tempered Clavier* are a particularly consistent corpus in their exploration of the 24 major and minor tonalities in 48 pairs of preludes and fugues. Fugues are often viewed as one of the highest forms of classical Western music, and therefore are often used in the teaching of music analysis and of composition (Bergerault 2011). Although its structure may appear formal, the fugue enables high levels of creativity. Many treatises have been written on fugues and on counterpoint in general, such as those by François-Joseph Fétis

(1825), André Gedalge (1901), Hugo Norden (1977), and Kent Kennan (1999). Some studies have included a complete musicological analysis of Bach's *Well-Tempered Clavier*, such as the books by Ebenezer Prout (1910) and Siglind Bruhn (1993).

The analysis of a fugue requires an understanding of its structure on both local and global scales. Fugues are thus good candidates for music information retrieval (MIR) research, stimulating the development of algorithms on symbolic scores. We believe that this research should go beyond generic problems such as pattern matching or inference or local key estimation. The research should include algorithms that predict actual “analytical elements” that will contribute to a large-scale analysis.

We describe here a system for automated fugue analysis, starting from a symbolic score that is separated into voices. We complete and extend methods to detect subjects and countersubjects, based on an analysis of repeating patterns, and methods to study episodes focusing on harmonic sequences (Giraud, Groult, and Levé 2012, 2013). Our aim is to provide algorithms with high precision rather than high recall. As all the content of a fugue is somewhat derived from the base patterns, what

Figure 1. Opening measures of J. S. Bach's Fugue in C minor (Fugue No. 2, BWV 847). Subjects (S), countersubjects (CS1 and

CS2), and first two episodes (E1 and E2). The notes ending the subjects, as described in the ground truth file, are circled, and the notes ending the

countersubjects are boxed. Inside the episodes, the inner brackets show the recurring patterns (three occurrences in E1, and two occurrences in E2).

The figure displays the opening measures of J.S. Bach's Fugue No. 2 in C minor, BWV 847, across three systems of musical notation. The first system shows the initial entry of the Subject (S) in the alto voice, followed by the First Counter-Subject (CS1) in the alto voice. The second system shows Episode 1 (E1, codetta) with the First Counter-Subject (CS1) in the soprano voice and the Second Counter-Subject (CS2) in the alto voice. The third system shows Episode 2 (E2) with the Subject (S) in the tenor voice and the First Counter-Subject (CS1) in the tenor voice. Various annotations include voice labels (soprano, alto, tenor), episode labels (Episode 1, Episode 2), and specific subject/counter-subject labels. Brackets indicate recurring patterns within the episodes.

is interesting is not to locate as many approximate occurrences as possible or to infer very short patterns, but to provide an analysis with some *semantics*: The occurrences of the patterns and of other analytical elements should be organized into a meaningful analysis.

The next sections provide background information on fugues and on the related work in MIR. We then present a ground-truth analysis of a corpus of 24 fugues by Bach and 12 fugues by Dmitri Shostakovich. We then explain our generic strategy for matching of repeated patterns with diatonic similarities, and discuss algorithms studying specific elements, including thematic patterns, harmonic sequences, cadences, and pedals. We also explain the combination of these different analyses in a global analysis. The last section discusses our implementation, describes an interface for visualization, and compares the results of these algorithms with the ground truth.

Elements of Fugues

The following components of fugue analysis are illustrated by the Fugue in C minor (Fugue No. 2, BWV 847) from the first book of Bach's *Well-Tempered Clavier*, which has a very regular construction. A fugue is structured as a set of *voices*, where each voice is a mostly monophonic sequence of notes. Fugue No. 2 consists of three voices. A fugue is built on a theme called a subject (S). The first three *entries* (or *occurrences*) of the subject in Fugue No. 2 are shown in Figure 1: The subject is stated in one voice (the alto), beginning on C, until the second voice enters (the soprano, measure 3). The subject is then stated in the second voice, this time transposed up a fifth. Meanwhile, the first voice continues with the first countersubject (CS1), which combines with the subject.

The fugue alternates between other entries of the subject and the countersubjects (eight instances

Figure 2. End of the Bach Fugue in C minor (Fugue No. 2, BWV 847). The concluding perfect authentic cadence (PAC, middle of measure 29) is followed by a bass pedal and a final exposition of the subject in the soprano voice.



of S, six instances of CS1, and five instances of the second countersubject, CS2), and developments on these patterns called *episodes* (E). Figure 1 shows the first two episodes of Fugue No. 2. The first episode is sometimes called *codetta*; the term “episode” can be applied to the ones after the exposition of all voices. The episodes can contain *cadential passages* that release tension. They often consist of *harmonic sequences*, which are passages where a pattern is consecutively repeated starting on a different pitch, possibly modulating from one tonality to another. The end of the fugue frequently includes a *stretto* (i.e., a succession of overlapping, often incomplete entries of S, in all voices). The final *cadence* is typically followed by a *bass pedal* (i.e., a note sustained over several harmonies) possibly with a last exposition of the subject (see Figure 2). Figure 3 shows an outline of the entire Fugue No. 2.

The subject and the countersubject form the foci of cognition when listening or playing to a fugue. As stated by Bruhn, “The perfect little musical entity we call subject is in fact at the origin of the fugue. . . . The subject is responsible for the feelings of density and relaxation in the fugue, and it is the main force in creating structure” (Bruhn 1993, p. 43). Therefore, most fugue analyses first consider these thematic patterns. On the other hand, the episodes induce variety and contrast. Messiaen (1956, p. 40) even stated: “Without constraining ourselves to making regular fugues, we shall keep the most essential parts of them: the episode and the stretto.” Indeed, fugal passages in Messiaen’s music are very far from the usual schema.

In regular fugues, it is the succession of the exposition and the episodic parts that builds the musical interest through contrasting phases of tension and release: “An episode can establish three

crucial relationships to its surroundings: It can link two subject statements by leading from one toward the next, it can be conclusive by resolving tension that was built up during a preceding subject statement, or it can represent a different register and serve as a color contrast” (Bruhn 1993, p. 50). This succession of exposition and episodic parts can be revealed by musicological analysis, and forms the basis of further aesthetic analysis.

Related Work: MIR and Fugue Analysis

Many generic MIR techniques can be applied to fugues, mostly in searches for repeating or musically salient elements (e.g., motifs, patterns, themes, or streams). For example, algorithms can be used to compute the similarity between monophonic sequences, such as the algorithms described by Mongeau and Sankoff (1990) and their extensions. One can also use methods for approximate pattern matching (Crawford, Iliopoulos, and Raman 1998; Clifford and Iliopoulos 2004), allowing a given number of mismatches.

Discovering patterns in music is an established field of MIR research (Jansen et al. 2014). Several studies have focused on finding maximal repeating patterns, thus discarding patterns that are sub-patterns of a larger pattern that occurs with the same frequency as the sub-patterns (Hsu, Liu, and Chen 1998; Liu, Hsu, and Chen 1999; Karydis, Nanopoulos, and Manolopoulos 2007). The repeating patterns can be searched using encodings more adapted to melody recognition, such as in the step-leap representation used by Cambouropoulos and colleagues (2005).


Figure 3. Analysis of Bach Fugue No. 2 in C minor from the first book of the Well-Tempered Clavier (BWV 847). Extract from the ground truth file (a), showing measures that contain entries of the subject (S), first countersubject (CS1) in soprano (S), alto (A), and tenor (T) voices. The position of each occurrence is given by the number of the measure in

which it occurs (3 is the first beat of measure 3). The actual start is, moreover, shifted from this logical position (one eighth for S, one sixteenth for CS1, except for some occurrences where there is an additional “start” keyword). This fugue also contains perfect and imperfect authentic cadences (PAC and IAC) on the tonic (i, I) and dominant (v). Uppercase is

used to indicate major keys and lowercase indicates minor.

Visualization of this ground truth file, and results of the proposed system for the same fugue (b). The algorithms retrieve all the entries of the subject (and, in Roman numerals, the scale degrees of the subject relative to the global key), occurrences of the “head” of the subject (small boxes

inside the episodes), all but one of the entries of the countersubject and second countersubject, harmonic sequences (1, 2, 3), three out of the four cadences and the final pedal. The cadences are denoted here by their root. In this example the algorithm has a perfect precision.



Other studies try to find musically significant elements, with algorithms considering the number of occurrences (Smith and Medina 2001), but also the melodic contour or other features (Meek and Birmingham 2003). Lartillot (2007) proposed a theory of “motivic pattern mining,” inferring an exhaustive list of repeating patterns, taking into account cognitive elements. Information about repetition can also be used, as in the local boundary detection model (cf. Cambouropoulos 2006). Many different features of theme extraction have also been discussed by Meek and Birmingham (2003).

Several techniques have been put forward to deal with polyphonic data. Meredith, Lemström, and Wiggins (2002; see also Meredith 2006a) have described a geometric encoding of notes (length, pitch) using techniques such as translation vectors for pattern discovery and matching. Conklin and Bergeron (2010) implemented a method to find short contrapuntal patterns that embed relations of consonance and dissonance between voices. Rafailidis et al. (2008) consider information about

melody, pitch, and rhythm to gather notes into stream segments, predicting phrasal boundaries but also some vertical groupings.

Conklin (2010a,b) focused on maximally general distinctive patterns, overrepresented in a corpus with respect to an “anticorpus.”

Many of these techniques may be used to give relevant pattern information in fugues. Any person or algorithm trying to do an actual analysis should, however, make wise choices of which elements to consider in undertaking a meaningful analysis. Using or extending some of these techniques, several MIR studies have already focused on fugues, attempting to achieve a musicological analysis of this repertoire. Weng and Chen (2005) built a tool to decide whether a piece is a fugue, but no details were given regarding the algorithm. Browles (2005) worked on fugue analysis, investigating several heuristics to aid in the selection of candidate fugue subjects, using algorithms for repeated pattern extraction developed by Hsu, Liu, and Chen (1998). These algorithms maximize the number of

occurrences of repeating patterns. The Web site www.earsense.org has also produced analyses of fugues by extracting sequences of some repeating patterns, but without precise formal analysis or precise bounds.

A Corpus of Fugues with Ground-Truth Analysis

Because evaluating algorithms is always challenging in MIR, there is the need for public data sets to evaluate and compare different algorithms. Previous works on fugue analysis were not evaluated against detailed ground-truth data. The evaluation developed by Weng and Chen (2005) was limited to determining whether a piece was a fugue or not. Browles' bachelor's thesis (2005) reported that the fugal subjects identified by her method were "missing or including an extra 1 to 4 notes." She discussed some individual results but did not describe any formalized ground truth.

We present here our corpus and the musicological sources, and argue why a specific ground truth for fugues is meaningful. The goal is to provide in a computer-readable format analytical elements that could be used in a musicological analysis of the fugue.

The Corpus and the Musicological Sources

We selected a corpus of the 24 fugues in the first book of Bach's *Well-Tempered Clavier* (BWV 846–869), and the first 12 fugues in Shostakovich's *24 Preludes and Fugues* (op. 87, 1952). Bach's *Well-Tempered Clavier* has been extensively studied, and systematic analyses of Bach's fugues have been published by Prout (1910), Tovey (1924), Keller (1965), and Bruhn (1993). All these studies discuss the S, CS1, and CS2 patterns and global characteristics of each fugue. Moreover, the detailed books by Bruhn list and discuss all entries of patterns. Charlier (2009) offers multi-thematic fugal analyses, some of which differ from more generally accepted ones. Shostakovich's fugues have received less attention than Bach's fugues. One systematic study has focused on their interpretation (Plutalov 2010), but it does not include a detailed formal analysis of the positions of

patterns. The Web site www.earsense.org provides some analytic elements of the fugues of both Bach and Shostakovich, including a discussion of their S, CS1, and CS2 patterns and the number of their entries (but not their positions).

Is a Ground Truth Meaningful?

There is not just one correct analysis of a given piece—musicologists often disagree, or they will at least advocate several different points of view. There may be no consensus among musicologists regarding even some basic analytical elements. The notions of "subject," "countersubject," or "cadence" are not so well defined and involve implicit or explicit musical knowledge: Somehow, if these notions were perfectly formalized, we could have exact algorithms giving perfect results. Moreover, these notions are more musicological (discussed in treatises) than empirical (tested with music cognition experiments).

In 8 of the 24 fugues by Bach, at least two sources disagreed on the end of the subject (see Figure 4). This ambiguity can be a part of the music, building an elaborate construction of the piece. On Fugue No. 9, in E major, Tovey (1924, p. 79) writes: "It is not worthwhile settling where the subject ends and where the countersubject begins." There is indeed a continuous flow from the subject to the countersubject.

We still argue that some formalized ground truth is meaningful. In 16 of the 24 same fugues, all our sources (except Charlier) agree exactly on the definition of the subjects. This consensus can be also found with other elements—the sources agree on the analytical elements on the very regular Fugue No. 2 depicted in Figure 3. The goal of such a ground truth is thus to formalize what is commonly perceived as analytical elements by traditional musicological analysis.

The Ground Truth

Prior to the present study, no computer-readable analysis has been performed on the fugues of either Bach or Shostakovich. We propose the first

Figure 4. The eight subjects in the first book of Bach's Well-Tempered Clavier where at least two sources disagree on the end of the subject. Circled notes show the proposed endings of the subjects. (Charlier has a motivic approach, resulting in shorter subjects.) Boxed notes show the ends found by our method.

05 Charlier All

07 Prout Bruhn Keller

09 Bruhn Prout Keller

10 Keller Prout Bruhn

11 Prout Keller Bruhn

18 Charlier All

19 Prout Bruhn Tovey Keller

22 Prout Bruhn Keller Bruhn

formalized ground-truth analysis of these fugues, which is usable in MIR evaluation and is built on the aforementioned sources and our own analysis. The ground-truth files (see Figure 3a) give the symbolic positions (measure number and position in measure) of the entries of the subject and the countersubject, as well as cadences and pedals. We also report slight modifications to S, CS1, and CS2, such as actual start, delayed resolutions, and other potentially useful information. Hence, we provide some analytic elements that could be part

of a musicological analysis. We do not claim to provide a full musicological analysis, however—these analytical elements are not the goal of the musicological analyses, but simply a part of these analyses. In particular, we do not analyze some high-level concepts such as the texture or the evolution of tension, nor the detailed transformation of the S, CS1, and CS2 patterns.


The ground-truth data should include analyses by different musicologists. But different analysis will focus on different aspects that are not easily

The current release of our data set (2015.01) includes the position of all complete subjects and countersubjects, as well as cadences and pedals—a total of more than 1,000 annotations. These files are available online at www.algomus.fr/datasets under open-source database licenses. In addition to evaluating fugue analysis, this data set might be useful in other situations besides the evaluation of fugue analysis, such as the evaluation of algorithms for pattern extraction or structure analysis.

A fugue contains many repeated patterns in either entries or episodes. Not all the instances of S, CS1, and CS2 are exact repetitions, however, and the patterns can be transposed or altered in various ways. Similarly, the patterns that form harmonic sequences in episodes are not always repeated exactly. Starting with the voice-separated data, we aim to detect such repeating patterns under a substitution function by considering the diatonic similarity of pitch intervals and by redefining the duration of all except the first and the last notes to take on the same value (the duration of the notes at the extremities of a pattern are altered more frequently).

A note x is described by a triplet (p, o, d) , where p is the pitch, o the onset, and d the duration. The pitch is an integer describing information that can be either diatonic (based on the succession of notes in a scale) or based on semitones. We consider ordered sequences of notes $x_1 \dots x_m$, that is, $x_1 = (p_1, o_1, d_1), \dots, x_m = (p_m, o_m, d_m)$, where $0 \leq o_1 \leq o_2 \leq \dots \leq o_m$ (see Figure 5). A sequence is *monophonic* if two notes never sound at the same onset, that is, for every i with $1 \leq i < m$,

In this example, onsets and durations are counted in sixteenth notes, and pitches and intervals are counted in semitones.



The similarity score between a pattern and the rest of the fugue can be computed via dynamic programming, using equations similar to those developed by Mongeau and Sankoff (1990). Let $x = x_1 \dots x_m$ be a pattern and $y = y_1 \dots y_n$ be a voice. When $a \leq m$, let $S(a, b)$ be the best number of matched intervals when aligning the start of the pattern $x_1 \dots x_a$ against $y_i \dots y_b$, a part of the voice y finishing at b , and let $S_f(m, b)$ be the best number of matched intervals when aligning the complete pattern $x_1 \dots x_m$ (a candidate pattern) against the same part $y_i \dots y_b$. The tables S and S_f are computed using the following dynamic programming equation:

The alignment then can be retrieved through backtracking in the dynamic programming table. The substitution functions δ and δ_f are defined as follows (the relation $\approx$ is a similarity relation on

Figure 6. Alignment between four similar patterns in Bach Fugue No. 5. Considering diatonic steps, the pitch intervals of the first pattern (m2) are $(\Delta x_2 \dots \Delta x_6) =$

$(-1, -1, -1, +1, -4)$. The pattern would be exactly matched in the diatonic transposition T, which does not occur in the fugue. In the actual occurrences of the pattern,

the first note (circled) is either fragmented (in m8) or substituted (in m5), and the length of the last note (circled) is not preserved. The pattern is further modified by an insertion

(boxed pitch B) that is better viewed as a fragmentation of the preceding note, D.



pitch intervals):

$$\delta((^{\Delta}p, o, d), (^{\Delta}p', o', d')) = \begin{cases} +1 & \text{if } ^{\Delta}p \approx ^{\Delta}p' \text{ and } d = d' \\ 0 & \text{if } ^{\Delta}p \not\approx ^{\Delta}p' \text{ and } d = d' \\ -\infty & \text{otherwise } (d \neq d'). \end{cases} \quad (2)$$

$$\delta_f((^{\Delta}p, o, d), (^{\Delta}p', o', d')) = \begin{cases} +1 & \text{if } ^{\Delta}p \approx ^{\Delta}p' \\ 0 & \text{otherwise } (^{\Delta}p \not\approx ^{\Delta}p'). \end{cases} \quad (3)$$

Note that δ checks pitch intervals and durations, whereas δ_f only considers intervals, ignoring the durations of the last notes. Neither of the durations of the first notes (x_1 and y_1) is checked, as the algorithm actually compares $^{\Delta}x_2 \dots ^{\Delta}x_a$ against $^{\Delta}y_2 \dots ^{\Delta}y_b$. This is because the first and the last notes of thematic patterns are transformed more frequently than the other notes, as in Figure 6.

These substitution functions are very conservative, allowing only substitution errors and strict length matching to retrieve as few false positives as possible while keeping a high recognition rate. For example, in the alignment of Figure 6, the pattern at measure 2 perfectly matches its diatonic transposition T: $S_f(m2, T) = 5$. As soon as the pattern is mutated, however, the pattern is not recognized: here $S_f(m2, m8) = S_f(m2, m5) = -\infty$. The equations can be extended to consider other editing operations. For example, allowing insertion


or deletion operations with a score of zero will give $S_f(m2, m8) = 4$ and $S_f(m2, m5) = 2$. Appropriate fragmentation operations with a score of 1 will raise the score $S_f(m2, m5)$ to 4. The comparison of S_f values with a threshold will decide whether the occurrence is kept or not. A threshold of 3 will allow here up to two interval mismatches. However, when the constraints are relaxed, because almost all the content of a fugue is somewhat derived from a subject or some countersubject, any part would match a part of the subject or another base pattern within a given threshold.

Many similarities $\approx$ can be selected, from a strict pitch equality to very relaxed “up/down” classes defining the melodic contour (Ghias et al. 1995). Some intermediary interval classes may be defined as “step/leap intervals” (Cambouropoulos et al. 2005) or “quantized partially overlapping intervals” (see Lemström and Laine 1998). The quantized partially overlapping interval model defines *short* intervals (from one to three semitones), *medium* intervals (from three to seven semitones), and *large* intervals (starting from six semitones). When the pitch names are available—or when they can be computed with pitch spelling algorithms, such as those developed by Meredith (2006b)—the similarity can be directly computed on *diatonic* pitches. A diatonic model is relevant for tonal music because it is sensible enough to allow mode changes while remaining specific—that is, a scale will always match only a scale (Orpen and Huron 1992; Cambouropoulos 1996; Hiraga 1997; Perttu 2000).

Figure 7. Relative start (s, s') and end (e) positions of countersubjects (1:CS1, 2:CS2) described in the ground truth file. These positions are computed as a number of notes

compared to the corresponding S start and end positions. The histograms show the distribution of these values on the 36 fugues of the corpus. Usually $s'_1 = 0$

(CS1 starts right after the subject) and $e_1 = e_2 = 0$ (CS1 and CS2 finish exactly at the end of the subject).



Algorithms for Fugue Analysis

We have developed several tools that give structural information about fugues. Using the adapted Mongeau-Sankoff algorithm and other discrete techniques, we aim to compute the end of the subject and then to infer CS1 and, possibly, CS2. Determining these patterns and their complete occurrences (including augmented or inverted ones) and their transpositions will help to shape the structure of the fugue. We identify harmonic sequences, cadences, and pedals. We also combine some of these elements into a global segmentation with statistical modeling. The result of these algorithms on Bach's Fugue No. 2 was depicted at the bottom of Figure 3 (in the block labeled "Computed analysis"). The output of our method is thus both the different elements and the prediction of the global segmentation.

Subject Identification

Finding the end of the subject is necessary for finding the countersubject and for building the complete structure of a fugue, including all the occurrences of these patterns. At the beginning of a fugue, the subject is first heard alone and is played by a single voice until the second voice enters. Fugue No. 6 from Shostakovich's op. 87 is a counterexample in which the first entry of the subject, in the bass voice, is doubled at the octave, although this can be viewed as a difference of texture.

The end of the subject is generally not exactly at the start of the second voice, however. For

example, in Bach's Fugue No. 2, according to the four musicological references of our ground truth, the first subject ends on an alto note E_b , which is two notes (two sixteenths) before the start of the soprano voice (i.e., the first circled note in Figure 1). Several parameters contribute to this perception. Metrically, the phrase ends on a strong beat; harmonically, the five preceding notes, $F G A_b G F$, suggest a dominant seventh chord, which resolves on the E_b , suggesting the C minor tonic. The subject ends with a succession of sixteenth notes with small intervals, whereas the following note (C) belongs to CS1, being separated from the E_b by a large leap and starting a descending scale of sixteenth notes. In the 36 fugues of our corpus, the end of the subject is between -8 and $+6$ notes (inclusive) relative to the start of the second subject (except for Shostakovich's Fugue No. 6 in B minor, -18 notes), and in the majority of cases, between -4 and $+1$ notes (inclusive, see s_1 in Figure 7).

Here, a simple algorithm, based only on similarities, is able to precisely detect most of the subject ends (Giraud, Groult, and Levé 2013). In the first voice, we test patterns that finish between eight notes before and six notes after the start of the second voice: Let z be the position of the note of the first voice heard before or at the start of the second voice, and $g \in [-8, 6]$. Each of these candidates, finishing at $z + g$, is matched against all the voices. This matching is done with the function described in the previous section (strict length matching except for the first and the last notes and only pitch interval substitution). The "best" one (that is, the candidate having the best total score on all its

occurrences) is then selected. The chosen subject is the one finishing at $z + g$ that maximizes the sum of the scores of its occurrences, that is, the sum of the $S_f(z + g, i)$ values for some position i , such that $S_f(z + g, i)$ is greater than or equal to a threshold, depending on the length of the pattern.

Moreover, a subject is more likely to end on a strong beat, and the pitch of the last note of the subject is, in most cases, a note of the tonic chord. These constraints could be encoded as additional scores in the matching process in order to improve the detection in some cases.

Countersubject Identification

The CS1 usually starts immediately after the subject (see s'_1 in Figure 7), except in some cases, such as in Bach's Fugue No. 16 in G minor, or in Shostakovich's Fugue No. 2 in A minor. CS2, when it exists, starts later, approximately at the start of the third entry of the subject (possibly after a codetta) and ends approximately at the end of this third subject (and at the second entry of CS1). Figure 7 shows the distribution of these relative positions in the corpus. In the majority of cases, countersubjects end at exactly the same position as the subjects (see boxed notes on Figure 1). This reinforces the perception of closure of the S/CS1/CS2 set, and helps to mark a transition between the exposition and the episodic material. To discover CS1, we test patterns starting just after the subject and finishing roughly after the end of the second subject (between two notes before and four notes after). We select the best candidate, again by choosing the z that maximizes the sum of $S_f(z + g, i)$ for all occurrences i of the candidate pattern. Note that a bad S detection can thus lead to a bad CS1 detection. The CS2 is even more difficult to find, because it can be shorter, and its starting position is difficult to predict. Moreover, musicologists do not always achieve consensus on the presence of CS2.

Detection of All Occurrences

Once the S, CS1, and CS2 patterns have been identified, it is possible to launch further dynamic

programming to detect more occurrences of these patterns. The very conservative substitution functions δ and δ_f could be relaxed to allow duration substitutions or, at least, fragmentations and consolidations that preserve the total duration. When we lower the global threshold, we find incomplete occurrences of S that are not proper entries. These relaxed constraints should not be used in the pattern identification phase because they are insufficient to detect endings. Moreover, we specifically look at the "head" of a subject by searching patterns limited to the first five notes of the subject (see Figure 3).

We add other constraints for the detection of CS1 and CS2, allowing a CS1 only when there is a concurrent S and a CS2 only when there is a CS1 in a "compatible" position. Finally, in some fugues, the subject appears inverted (or "upside down", i.e., all intervals are mirrored) or augmented (e.g., all durations are doubled). When the subject is known, the same matching algorithm retrieves any inversions and augmentations of the subject.

Subject Scale Degrees and Local Tonalities

After an exposition using the tonic and dominant (I and V degrees), the fugue often moves to relative or more distant tonalities, returning to the tonic for the conclusion. Determining these local tonalities is a part of a complete analysis. Such tonalities could be estimated by using the classical algorithms described in the works of Krumhansl and Kessler (1982) or Temperley (1999), and their extensions, or using the more recent findings of Madsen and Widmer (2007) or Robine, Rocher, and Hanna (2008). These algorithms may, however, incorrectly assign tonalities in rapidly modulating passages, such as some episodes.

We do not use any key finding but instead use a much simpler method, taking advantage of subject detection and not trying to study the tonality in episodes. We focus on the subjects (which are points in which the local tonality is more stable), and determine the transposition interval between the current subject and the first subject, which is at the tonic, I. For this, we compute the average pitch, in semitones, of each subject entry. The scale degrees of

Figure 8. Bach Fugue No. 22 in B $\flat$ minor. The subject is stated in the soprano in measures 55–57. The line of parallel thirds between soprano and alto starts a pattern used for the

following episode. A partial harmonic sequence is detected starting from measure 56, lasting for two phrases of two measures (eight quarter notes) each—this sequence is not

exact, as the end of the pattern is varied. Another partial harmonic sequence with 2 $\times$ 1 measure is detected later.

the subjects are then computed relative to the global key as established for the first subject. When the subject is transposed from the tonic to the dominant, some slight modifications in the subject can happen (known as a *tonal answer*), usually in the very first notes of the subject (for instance, the dominant pitch may be transposed to the tonic pitch instead of the supertonic). These modifications are compensated for by the other notes in the average pitch, so the computation of the scale degree is still correct. This does not allow the detection of changing tonalities inside a subject, however. Many subjects indeed start from the tonic and modulate towards the dominant. Knowing the degrees of the greater part of the subjects gives insight, nevertheless, into the tonal path of the complete piece.

For example, this tonal path can be seen in Bach's Fugues No. 2 (see Figure 3) and No. 18 (G $\sharp$ minor, discussed subsequently); also note that scale degrees are denoted in roman numerals at the bottom of the diagrams. Whereas the initial exposition and the ending of the fugue alternate between the tonic (I) and either the subdominant or the dominant (IV or V), the middle sections reach the relative key (III or VI) or more distant keys.

Harmonic Sequences

In fugues, harmonic sequences are almost exclusively found in the episodic parts. Detecting them helps to reveal the fugue's design. In each voice, we look for consecutive repeating patterns—that is, two consecutive occurrences pat_1 and pat_2 of the same pattern, and possibly three (pat_1 , pat_2 , pat_3). We therefore use the substitution function described in the previous section, which requires

that $S(pat_1, pat_i)$ is above a given threshold for each occurrence pat_i .

We try to detect partial harmonic sequences starting at every note occurring at or immediately after each beat of the score. A sequence is reported when patterns in at least two different voices are repeated consecutively with the same transposition (Giraud, Groult, and Levé 2012). Several periods (pattern length) are tested: one and two beats (depending on the meter), one and two measures. Because every beat is tested, this procedure finds several overlapping sequences, with the same period, as soon as there are strictly more than two occurrences. Such overlapping sequences are merged into a single sequence.

This method correctly retrieves most of the harmonic sequences. Near the beginning of Fugue No. 2 (see Figure 3), the two harmonic sequences are detected. The first one, during the codetta, is shortened because of the shifted start of the two voices, however. The harmonic sequences are good markers of episodes. A frequent situation, however, is that the harmonic sequence begins at the end of a thematic pattern, continuing in episode (see Figure 8).

Cadences and Pedals

Cadences and pedals mark structural transitions in the fugue (ends of expositions, ends of episodes, and coda with a pedal). Except for some particular cases (see Figure 9), the pedals are found at the very end of the fugues, as in Fugue No. 2 (see Figure 2). We label as “bass pedal” every bass note that lasts strictly more than four beats (for binary meters) or six beats (for ternary meters).

Figure 9. Unusual cases of pedals. In the middle of Fugue No. 11 in F major (a), there is a pedal, which has been correctly detected, in the dominant of the relative key. This pedal in Fugue No. 18 in G# minor (b) starts in the tenor and then continues in the bass. It could be detected by geometrical methods.

Figure 9 consists of two musical staves, (a) and (b), illustrating unusual cases of pedals. Staff (a) shows measures 36 and 37 of Fugue No. 11 in F major. It features a soprano (S) and alto (S) voice part with a pedal in the bass. Staff (b) shows measures 37 and 38 of Fugue No. 18 in G# minor. It features a pedal starting in the tenor and continuing in the bass.

Each time the bass moves up a perfect fourth or down a perfect fifth, we try to identify a cadence V-I in root position if the following conditions are met. On the beat, the notes of the chord must be part of a major, minor, or suspended fourth chord with the bass tonic as root. The uppermost note must be either the tonic (perfect authentic cadence, PAC) or a major or minor third (root position imperfect authentic cadence rIAC). Note that the arrival chord is often not a full perfect chord, and sometimes it can even be limited to the single tonic note, such as in the PAC of Figure 2. Before the beat, at some point during the bass note, we must find a dominant chord (notes included in a dominant seventh, with at least the leading tone). Because of passing tones, this point is often not the very last note before the strong beat.

This method allows us to find about 82 percent of the PACs in the ground truth. Most of the false positives are cases where these conditions are met but the musical phrase continues (see Figure 10).

Global Structure

We use a hidden Markov model (HMM) to combine different elements of analysis. An HMM consists of a sequence of *states* with *transition probabilities*

and *emission probabilities*. At each computation step, the HMM can change its state according to a transition. The HMM emits symbols, which are the only output of the system; the sequence of the corresponding states is hidden. Given an observed sequence of symbols, the Viterbi algorithm finds the most likely sequence of states by dynamic programming. Markov models have several applications in MIR. For example, they have been used to predict musical labels on segmentations (Paulus and Klapuri 2009).

In fugue analysis, the goal of this model is to discover the sequence of states that structures the fugue (exposition, codetta, further expositions, and further episodes, see Figure 11). The observed symbols are the primary elements of the analysis (thematic pattern entries, harmonic sequences), instead of the raw notes. The symbols are for each quarter note of the score. Examples of the HMM's output can be seen at the bottom of Figures 3, 12, 13, and 14.

Implementation and Visualization of the Analysis

All the algorithms described were implemented in Python. The music21 framework (Cuthbert and Ariza 2010, version 2.0.0) was used for input (parsing

Figure 10. Examples of discovered cadences. Fugue No. 18 in G# minor, measures 8–11 (a). There is an interval of a perfect fifth at the end of the subject, leading to spurious cadences when the subject is stated in the bass. The algorithm detects three

cadences here, at the first beat of each measure. The first one (first beat of measure 9), can be viewed as a real PAC finishing the exposition of thematic patterns. The middle cadence (first beat of measure 10) is certainly a false positive. The last one

(first beat of measure 11) can be viewed as a real PAC finishing the harmonic sequence. This is reinforced by the change of texture in measure 11, with a new S entry at the tenor on the second beat. Fugue No. 20 in A minor, measures 46–48 (b). Our

algorithm finds three consecutive cadences, but only the one ending on the first beat of measure 48 is a real PAC, the notes on the third beat of measures 46 and 48 being parts of larger phrases.




Figure 11. Sketch of the hidden Markov model used for fugues. The transition probabilities were manually selected to favor some stability in the

states, as the sections usually span at least one measure. The first exposition and the codetta states are used only during the exposition. The

emissions shown with dotted arrows are rare, with small probabilities (harmonic sequences during exposition sections).




Figure 12. Ground truth and results for Bach Fugue No. 18 in G# minor. Incomplete CS2s in the ground truth are denoted by “inc.” The computed

analysis includes several spurious cadences because there is an interval of a perfect fifth at the end of the subject (see Figure 10).

Figure 13. Ground truth and results for Shostakovich Fugue No. 8 in F# minor.




Figure 12




Figure 13

of MIDI files and Humdrum ^{**1}kern-format files). For performance reasons, however, our algorithms were based on custom objects that were “lighter” than the full-featured Stream and Note objects of music21. We implemented output as text, as scalable vector graphics (see Figures 3, 12, 13, and 14), and as HTML5 (www.algomus.fr/fugues). The Web output allows for switching between the ground truth and the computed analysis. It displays the score extracts of the patterns with the music21j javascript library (<https://github.com/cuthbertLab/music21j>).

Evaluation and Discussion

We ran the algorithms on the 36 fugues in our corpus. Bach’s fugues were taken from Humdrum ^{**1}kern files encoded by David Huron, which are available for academic purposes at KernScores (<http://kern.humdrum.org>). Shostakovich’s fugues were taken from MIDI files encoded by José Oscar de Almeida Marques (2002).

We compared all our results with the ground-truth file (version 2015.01). Table 1 shows the statistics of the results. All the algorithms, except

Figure 14. Ground truth and results for Bach Fugue No. 4 in C# minor, a triple fugue. The ground truth describes three subjects (S, S2, S3), some of which are incomplete (inc) or varied (var). The first part is

almost perfectly analyzed. The second fugue, starting on measure 35, is not analyzed. Analyzing this part would require the identification of new thematic patterns.




Table 1. Algorithm Predictions and Ground Truth

J. S. Bach, First Book of Well-Tempered Clavier

	ℓ	TP	FP	Recall	Precision
S	18/24	266/291	31	91.4	89.6
CS1	8/13	54/107	4	50.5	93.1
CS2	1/3	9/33	1	27.2	90.0
Sinv		21/24	2	87.5	91.3
Saug		3/3	0	100.0	100.0
PAC		57/69	13	82.6	81.4
rIAC		16/24	43	66.7	27.1
Pedal		15/20	10	75.0	60.0

D. Shostakovich, Op. 87, Fugues Nos. 1–12

	ℓ	TP	FP	Recall	Precision
S	5/12	115/172	14	66.9	89.1
CS1	2/5	42/99	1	42.4	97.7
CS2	2/4	19/51	0	37.3	100.0
Sinv		4/4	0	100.0	100.0

Comparison of the predictions of the algorithms with the ground truth file (version 2015.01), with the proportion of correct lengths found among the detected patterns (ℓ), the number of true positives (TP), false positives (FP), and the corresponding recall and precision. These values are calculated for the identification of fugue subject (S), countersubjects (CS1 and CS2), inverted and augmented subjects (Sinv and Saug), cadences (perfect authentic, PAC, and root-position imperfect authentic, rIAC), and pedal points.

for cadence identification, were found to have few false positives. In both the Bach and Shostakovich

corpora, the precision of pattern detection is more than 85 percent for S, CS1, and CS2. For example, false positive subjects were found in only three of Bach's fugues and in four of the fugues by Shostakovich. Most of the time, these seemingly erroneous occurrences are in fact relevant to the analysis, because they detect incomplete thematic patterns. Table 2 shows individual results for all the studied fugues. For each fugue, these results include a subjective evaluation of the quality of the output.

The algorithms were successful in finding the structure of several fugues, even if some elements were missing. For example, the relatively conventional Bach Fugue No. 2 was perfectly analyzed (see Figure 3), except the last CS1, which included a change of voice (a case that we did not include in the present study). The harmonic sequences were found in all the episodes, and the Markov model gave a perfect segmentation of the fugue. In Bach's Fugue No. 18, almost all entries of S, CS1, and CS2 were identified, as well as harmonic sequences (see Figure 12). The results of the automatic analysis of Shostakovich's Fugue No. 8 in F# minor were also very good (see Figure 13). One missing subject is an entry split between several voices (in this encoding). The two false positive subjects at measures 108 and 109 are relevant as long, incomplete occurrences of the subject.

In some fugues, the algorithms managed to give hints, but failed to provide a complete analysis of the important parts. Finally, the analysis of some fugues

Table 2. Results of Computational Fugue Analysis

Bach, First Book of *Well-Tempered Clavier*

<i>Fugue</i>	<i>S + Sinv/Saug</i>			<i>CS1</i>			<i>CS2</i>			<i>Qual.</i>	<i>Notes</i>
	ℓ	<i>TP</i>	<i>FP</i>	ℓ	<i>TP</i>	<i>FP</i>	ℓ	<i>TP</i>	<i>FP</i>		
1: CM	+	22/22								+	Perfect S detection, some cadences undetected
2: Cm	+	8/8		+	5/6		+	4/5		+	Almost perfect (see Figure 3)
3: C♯M	+	12/12		–	7/10		–	3/5	1	+	Good structure even if CS1 is too short and only the tail of CS2 is found
4: C♯m	+	13/23		+	2/2					–	Triple fugue, only the first fugue subject is retrieved (see Figure 14)
5: DM	–	11/11	25	–	0/9	4		0/6		–	Bad S detection (alternative S given by Charlier, but hides possible CS1)
6: Dm	+	11/12 + 3/5		–	2/3					=	Good structure, some S/Sinv and incomplete CS1 undetected
7: EbM	+	9/9			0/7					+	Good structure, almost perfect S, CS1 undetected, spurious cadences
8: D♯m	–	15/16 + 10/10	3 + 2							=	Good S/Sinv/Saug (even if bad S length), some spurious cadences
9: EM	–	10/12			0/3					–	Bad S detection (detected S overlaps true CS1)
10: Em	+	8/8			7/7					+	Good structure, almost perfect
11: FM	+	10/13		+	2/5					=	Some varied/fragmented occurrences of S are detected as heads of S
12: Fm	–	10/10		–	5/7			0/5		=	Good structure, spurious sequences, CS2/CS3 undetected
13: F♯M	–	7/8		+	2/5			0/3		+	Good structure, some spurious cadences, some patterns undetected
14: F♯m	+	6/7 + 2/2		+	4/6					+	Good structure, good S/Sinv/CS1
15: GM	+	4/4 + 2/3			0/1					–	Good S, but bad incomplete S detection
16: Gm	+	15/16		+	3/12					+	Good structure, almost perfect S, incomplete CS1 undetected
17: AbM	+	15/15								+	Almost perfect structure, even if some cadences badly detected
18: G♯m	+	12/12		+	5/8		–	2/3		+	Good structure, some spurious cadences (see Figure 12)
19: AM	+	9/10	1		0/2					=	CS1 occurs late and cadences undetected, FP S is varied S
20: Am	+	14/14 + 5/5			0/3					=	Good structure, good S/Sinv, many incomplete S
21: BbM	+	8/8		+	7/7			0/6		+	Good structure, even if CS2 undetected

Table 2. *Continued.*

Bach, First Book of Well-Tempered Clavier

<i>Fugue</i>	<i>S + Sinv/Saug</i>			<i>CS1</i>			<i>CS2</i>			<i>Qual.</i>	<i>Notes</i>
	ℓ	<i>TP</i>	<i>FP</i>	ℓ	<i>TP</i>	<i>FP</i>	ℓ	<i>TP</i>	<i>FP</i>		
22: Bbm	+	14/18								+	Good structure, some S and incomplete S undetected
23: BM	+	10/10 + 2/2		–	3/4					+	Good structure, Sinv detected, even if bad S length (shifted by one note)
24: Bm	–	13/13								=	Perfect S (but complex CS with several parts undetected)

Shostakovich, Op. 87, Fugues Nos. 1–12

<i>Fugue</i>	<i>S + Sinv/Saug</i>			<i>CS1</i>			<i>CS2</i>			<i>Qual.</i>	<i>Notes</i>
	ℓ	<i>TP</i>	<i>FP</i>	ℓ	<i>TP</i>	<i>FP</i>	ℓ	<i>TP</i>	<i>FP</i>		
1: CM	+	11/12		–	7/7		–	5/6		+	Good structure, almost perfect S/CS1/CS2
2: Am	–	9/10			0/4			0/3		=	Good S, but CS1/CS2 undetected
3: GM	–	7/13	2		0/7			0/4		–	Many S undetected, CS1/CS2 undetected
4: Em	–	7/12	2		0/7			0/4		–	Double fugue (only the first subject is partially retrieved, but bad S length)
5: DM	+	9/16			7/8	1		0/6		=	Some occurrences of S detected as head of S, bad length of CS1, CS2 undetected
6: Bm	–	0/22			0/12			0/6		–	Bad S length and CS1 undetected (CS1 too far from end of S)
7: AM	+	9/13		+	7/9		+	3/4		+	Good structure, almost perfect S/CS1/CS2, except changes of voice
8: F#m	–	8/10	2	–	2/7		–	2/3		+	Good structure, good S/CS1/CS, S FP are incomplete S (see Figure 13)
9: EM	+	11/11 + 4/4		–	5/5					+	Good structure, good S/Sinv occurrences (even if bad S length)
10: C#m	+	21/23		+	14/14		+	9/12		+	Good structure, good S/CS1/CS2
11: BM	–	9/14			0/7			0/3		=	CS1/CS2 undetected
12: G#m	–	14/16	8		0/12					=	Detects Sinc as S, CS1 undetected

Results from our corpus of 36 fugues, compared to the ground truth file (version 2015.01). For subject (S, Sinv, Saug) and countersubject (CS1 and CS2) patterns, the columns list good duration detection (ℓ), ratio of true positives (TP), and number of false positives (FP). A pattern occurrence is denoted as TP when it starts within two quarter notes of an occurrence of the ground truth. The column headed “Qual.” lists the subjective assessment of the quality of the complete system in sketching the global analysis of the fugue: good (+), acceptable (=) or bad (–). This assessment considers both isolated elements (subjects and countersubjects, sequences, cadences, and pedals) and the prediction of the structure.

yielded bad results that were almost unusable, especially when the subject detection failed. Some noteworthy cases were the double and triple fugues, where additional subjects entered with their own

expositions in the middle of the fugue, as in Bach’s Fugue No. 4 (C# minor; Figure 14). Because the detection is done at the beginning of the fugue, our system was not able to analyze these parts.

Conclusion

We presented a system for the complete analysis of fugues, starting from a symbolic score that was already separated into voices. Various MIR techniques were used to consider pattern repetition, harmonic analysis, and phrasing. Some of these techniques could also be used in the analysis of other genres.

We focused on algorithms with very high precision. S, CS1, and CS2 patterns were detected with false positive rates of approximately 10 percent. Finally, the use of the HMMs allowed sketching of a global structure, providing an acceptable or good structural analysis of about one half of our corpus. To evaluate our results, we published a corpus of 36 manually annotated fugues with more than 1,000 annotations, as well as a web page for interactive visualization.

We recommend further studies to improve the analysis of fugues. The tools used for the analysis could be improved for better recall, still keeping the high precision. For individual analytical elements, probabilistic models could be used instead of binary choices using thresholds (Temperley 2007). The Markov model could take into account more elements and predict more sections. Machine learning could improve the thresholds and weights of these models, but strategies have to be designed to address the problem of overfitting, a concern for data sets as small as these are prone. Future studies may include work on polyphonic data. The current system works on voice-separated data. Using plain MIDI files, one could use voice-separating algorithms or work directly on detection of patterns, matches, and cadences.

The algorithms could also be tested on other Baroque, classical, or Romantic fugues with the consideration of some practical limitations, such as the availability of voice-separated files and ground-truth data. As long as the fugues maintain a strict structure with a clear subject exposition, we are confident that our set of algorithms would yield good results. A further challenge is the successful analysis of more complex forms and the detection and analysis of fugue passages (*fugato*) in larger works.

References

- Bergerault, A. 2011. "L'enseignement du contrepoint et de la fugue au Conservatoire de Paris (1858–1905)." *Transposition: Musique et sciences sociales* 1. Available online at transposition.revues.org/418. Accessed December 2014.
- Browles, L. 2005. "Creating a Tool to Analyse Contrapuntal Music." Bachelor's thesis, Bristol University, Bristol, UK.
- Bruhn, S. 1993. *J. S. Bach's Well-Tempered Clavier: In-Depth Analysis and Interpretation*. Hong Kong: Mainer.
- Cambouropoulos, E. 1996. "A General Pitch Interval Representation: Theory and Applications." *Journal of New Music Research* 25(3):231–251.
- Cambouropoulos, E. 2006. "Musical Parallelism and Melodic Segmentation." *Music Perception* 23(3):249–268.
- Cambouropoulos, E., et al. 2005. "A Pattern Extraction Algorithm for Abstract Melodic Representations that Allow Partial Overlapping of Intervallic Categories." In *Proceedings of the International Conference on Music Information Retrieval*, pp. 167–174.
- Charlier, C. 2009. *Pour une lecture alternative du Clavier bien tempéré*. Jacquart.
- Clifford, R., and C. S. Iliopoulos. 2004. "Approximate String Matching for Music Analysis." *Soft Computing* 8(9):597–603.
- Conklin, D. 2010a. "Discovery of Distinctive Patterns in Music." *Intelligent Data Analysis* 14(5):547–554.
- Conklin, D. 2010b. "Distinctive Patterns in the First Movement of Brahms' String Quartet in C Minor." *Journal of Mathematics and Music* 4(2):85–92.
- Conklin, D., and M. Bergeron. 2010. "Discovery of Contrapuntal Patterns." In *Proceedings of the International Society for Music Information Retrieval Conference*, pp. 201–206.
- Crawford, T., C. S. Iliopoulos, and R. Raman. 1998. "String Matching Techniques for Musical Similarity and Melodic Recognition." *Computing in Musicology* 11:71–100.
- Cuthbert, M. S., and C. Ariza. 2010. "music21: A Toolkit for Computer-Aided Musicology and Symbolic Music Data." In *Proceedings of the International Conference on Music Information Retrieval*, pp. 637–642.
- Fétis, F.-J. 1825. *Traité du contrepoint et de la fugue*. Paris: Troupenas.
- Gedalgé, A. 1901. *Traité de la fugue*. Paris: Enoch.
- Ghias, A., et al. 1995. "Query by Humming: Musical Information Retrieval in an Audio Database."

- In *Proceedings of the ACM International Conference on Multimedia*, pp. 231–236.
- Giraud, M., R. Groult, and F. Levé. 2012. "Detecting Episodes with Harmonic Sequences for Fugue Analysis." In *Proceedings of the International Conference on Music Information Retrieval*, pp. 457–462.
- Giraud, M., R. Groult, and F. Levé. 2013. "Subject and Counter-Subject Detection for Analysis of the Well-Tempered Clavier Fugues." In M. Aramaki et al., eds. *From Sounds to Music and Emotions*. Berlin: Springer, pp. 422–438.
- Hiraga, Y. 1997. "Structural Recognition of Music by Pattern Matching." In *Proceedings of the International Computer Music Conference*, pp. 426–429.
- Hsu, J. L., C. C. Liu, and A. Chen. 1998. "Efficient Repeating Pattern Finding in Music Databases." In *Proceedings of the International Conference on Information and Knowledge Management*, pp. 281–288.
- Jansen, B., et al. 2014. "Discovering Repeated Patterns in Music: State of Knowledge, Challenges, Perspectives." In M. Aramaki et al., eds. *Sound, Music, and Motion*. Berlin: Springer, pp. 225–240.
- Karydis, I., A. Nanopoulos, and Y. Manolopoulos. 2007. "Finding Maximum-Length Repeating Patterns in Music Databases." *Multimedia Tools Applications* 32:49–71.
- Keller, H. 1965. *Das Wohltemperierte Klavier von Johann Sebastian Bach*. Kassel, Germany: Bärenreiter.
- Kennan, K. 1999. *Counterpoint*, 4th edition. Upper Saddle River, New Jersey: Prentice Hall.
- Krumhansl, C. L., and E. J. Kessler. 1982. "Tracing the Dynamic Changes in Perceived Tonal Organisation in a Spatial Representation of Musical Keys." *Psychological Review* 89(2):334–368.
- Lartillot, O. 2007. "Motivic Pattern Extraction in Symbolic Domain." In J. Shen, J. Shepherd, B. Cui, and L. Liu, eds. *Intelligent Music Information Systems: Tools and Methodologies*, pp. 236–260. doi:10.4018/978-1-59904-663-1.ch011
- Lemström, K., and P. Laine. 1998. "Musical Information Retrieval Using Musical Parameters." In *Proceedings of the International Computer Music Conference*, pp. 341–348.
- Liu, C.-C., J.-L. Hsu, and A. L. Chen. 1999. "Efficient Theme and Non-Trivial Repeating Pattern Discovering in Music Databases." In *Proceedings of the International Conference on Data Engineering*, pp. 14–21.
- Madsen, S. T., and G. Widmer. 2007. "Key-Finding with Interval Profiles." In *Proceedings of the International Computer Music Conference*, vol. 2, pp. 212–215.
- de Almeida Marques, J. O. 2002. "The Shostakovich Opus 87 Page." Available online at www.unicamp.br/~jmarques/mus/opus87. Accessed December 2014.
- Meek, C., and W. P. Birmingham. 2003. "Automatic Thematic Extractor." *Journal of Intelligent Information Systems* 21(1):9–33.
- Meredith, D. 2006a. "Point-Set Algorithms for Pattern Discovery and Pattern Matching in Music." In *Dagstuhl Seminar Proceedings: Content-Based Retrieval*. Available online at drops.dagstuhl.de/opus/volltexte/2006/652. Accessed December 2014.
- Meredith, D. 2006b. "The ps13 Pitch Spelling Algorithm." *Journal of New Music Research* 35(2):121–159.
- Meredith, D., K. Lemström, and G. A. Wiggins. 2002. "Algorithms for Discovering Repeated Patterns in Multidimensional Representations of Polyphonic Music." *Journal of New Music Research* 31(4):321–345.
- Messiaen, O. 1956. *The Technique of My Musical Language*. J. Satterfield, trans. Paris: Leduc.
- Mongeau, M., and D. Sankoff. 1990. "Comparison of Musical Sequences." *Computer and the Humanities* 24:161–175.
- Norden, H. 1977. *Foundation Studies in Fugue*. Carlsbad, California: Crescendo.
- Orpen, K. S., and D. Huron. 1992. "Measurement of Similarity in Music: A Quantitative Approach for Non-Parametric Representations." *Computers in Music Research* 4:1–44.
- Paulus, J., and A. Klapuri. 2009. "Labelling the Structural Parts of a Music Piece with Markov Models." In S. Ystad, R. Kronland-Martinet, and K. Jensen, eds. *Computer Music Modeling and Retrieval: Genesis of Meaning in Sound and Music*. Berlin: Springer, pp. 166–176.
- Perttu, S. 2000. "Combinatorial Pattern Matching in Musical Sequences." MSc thesis, University of Helsinki.
- Plutalov, D. V. 2010. "Dmitry Shostakovich's Twenty-Four Preludes and Fugues Op. 87." PhD dissertation, University of Nebraska.
- Prout, E. 1910. *Analysis of J.S. Bach's Forty-Eight Fugues: Das Wohltemperierte Clavier*. London: Ashdown.
- Rafailidis, D., et al. 2008. "Detection of Stream Segments in Symbolic Musical Data." In *Proceedings of the International Conference on Music Information Retrieval*, pp. 83–88.
- Robine, M., T. Rocher, and P. Hanna. 2008. "Improvements of Key-Finding Methods." In *Proceedings of the International Computer Music Conference*, pp. 281–284.

-
- Smith, L., and R. Medina. 2001. "Discovering Themes by Exact Pattern Matching." In *Proceedings of the International Conference on Music Information Retrieval*, pp. 31–32.
- Temperley, D. 1999. "What's Key for Key? The Krumhansl-Schmuckler Key-Finding Algorithm Reconsidered." *Music Perception* 17(1):65–100.
- Temperley, D. 2007. *Music and Probability*. Cambridge, Massachusetts: MIT Press.
- Tovey, D. F. 1924. *J.-S. Bach: Forty-Eight Preludes and Fugues (Das Wohltemperierte Clavier)*. London: Associated Board of the Royal Schools of Music.
- Weng, P.-H., and A. L. P. Chen. 2005. "Automatic Musical Form Analysis." In *Proceedings of the International Conference on Digital Archive Technologies*. Available online at nccur.lib.nccu.edu.tw/handle/140.119/23677. Accessed January 2015.