

Interactive data mining & machine learning techniques for musicology

Knut Hartmann, Daniel Büchner, Axel Berndt, Andreas Nürnberger

Department of Computer Science, University of Magdeburg, Germany
{Knut.Hartmann, Axel.Berndt}@isg.cs.uni-magdeburg.de
nuernb@iws.cs.uni-magdeburg.de

Carsten Lange

Zentrum für Telemann-Pflege und -Forschung Magdeburg, Germany
Carsten.Lange@tz.magdeburg.de

In: K. Maimets-Volk, R. Parncutt, M. Marin & J. Ross (Eds.)
Proceedings of the third Conference on Interdisciplinary Musicology (CIM07)
Tallinn, Estonia, 15-19 August 2007, <http://www.gewi.uni-graz.at/cim07/>

Background in history of music. The German Baroque music composer Georg Philipp Telemann was one of the most influential composers of his time with a large repertoire of musical styles. The reflection of musical influences of Telemann on other composers and the detection of musical influences in Telemann's *œuvre* are interesting research questions. Moreover, Telemann's approach to compositional problems, especially the treatment of the libretto in his vocal work, has been intensively studied during the last years. All these open musicological questions raise the demand for an extensive analysis of scores. Considering the pure amount of Telemann's *œuvre* and the multitude of corpus analysis tasks, the contrastive analysis of domain experts should be supported by powerful and adaptable software tools.

Background in computing, mathematics and statistics. Numerous algorithms were proposed to detect important musical features (melody, rhythm and harmony) with data mining and machine learning techniques in large corpora of scores. Moreover, the detection of melodic features often considers psychological theories on the reception of melodies. But these automatic feature extraction techniques need an unambiguous and complete problem description. A correct interpretation of musical phenomena, however, requires an in-depth knowledge of musical styles and compositional rules within a specific period of time. Therefore, musical experts should be able (i) to construct their corpus to be analyzed, (ii) to define queries which combine melodic, rhythmic, harmonic and lyrical aspects and (iii) to correct and augment search results.

Aims. The Telemann Research Center hosts a large archive of partly unpublished and undated scores of Telemann and a musicological library dedicated to the work of this composer. In a long-time effort, this corpus has been partly transformed into an outdated symbolic representation. These scores had to be converted into a more flexible musical notation (MusicXML). Moreover, flexible data mining algorithms had to be developed in order to extract those musical features which are interesting for a musicological analysis. Finally, flexible and adaptable graphical interfaces should support a contrastive analysis of large musical corpora. According to the main focus of this conference 'singing', the paper focuses on problems related to the musical realization of the libretti (for instance, the musical treatment of phrase endings in the vocal work of Telemann) and the detection of influences of ethnical and regional songs on Telemann's work.

The German Baroque music composer Georg Philipp Telemann (March 14, 1681 - June 25, 1767) wrote over 3,000 compositions. Sacred music constitutes the biggest part of his *œuvre* comprising over 1,700 cantatas, 16 masses, 23 psalms, 40 passion cantatas and 6 oratorios. For many compositions the dates of origin and even the authorship are unknown or at least questionable. Musicological investigations try to find answers by extensive manual corpus analyses. Algorithmic music analysis and data mining techniques are prospective tools to facilitate this work. We will describe an application scenario where these techniques proved to be of value as efficient time-saving tools in the research on Telemann's sacred music.

Problem Statement

Music analysis belongs to the armamentarium of musicologists. It is essential for a deeper understanding of the music of past epochs. Moreover, it provides invaluable insights into the development of today's music culture and listening consuetude. Musical styles and stylistic devices that are omnipresent and generally taken for granted in today's (e.g., western) music had to develop and establish in the compositional practice and acceptance of the audience—a process of centuries of an ongoing musical change.

Of special musicological interest are those composers that lived and worked in transitional periods, because their music reflects this stylistic change. One of them is Georg

Philipp Telemann. As one of the most influential composers of his time, his *œuvre* reflects the musical development from the baroque to the classical period of European music. But the big corpus of his compositions is only partially utilizable for further musicological investigations. Lots of arguable dating and authorship assumptions have to be clarified to provide a trustworthy solid scientific basis.

Telemann's cantatas may be taken as an example. It is supposed that some which are currently ascribed to J.S. Bach are composed by Telemann (Dürr, 1952). A contrastive analysis of stylistic aspects (see Fleischhauer, 1963; Hirschmann, 1992 and Reipsch, 1998b) can provide clues for a new evaluation of authenticity and attribution. Considering the correlation of compositional means to the text, Telemann's musical affect treatment can be a starting point for investigations (Lange, 2004).

The fact that Telemann wrote his cantatas in bigger consistent complexes each covering an "ideal" year of 72 Sundays can be of further help (Reipsch, 1998a). The trial of an attribution within such a complex can provide further hints.

The analysis of stylistic aspects and the mutual influence of contemporary composers and vogues furthermore give clues for the dating of origin. From Telemann it is known that he assimilated lots of cultural and personal influences from his journeys and correspondences in his own style. This is documented in his reports and letters (e.g., Telemann, 1767). Together with further contemporary sources the date of origin of particular compositions can be limited.

It is documented that Telemann got in touch with Hanacic style elements during a journey to Poland. Thus, the Lombardic rhythm, which was only detected in Telemann's music during a limited creative period by the musicologist Wolf Hobohm, can be used as an indication for further undated pieces (Hobohm, 1983). Therefore, of course, it should first be verified in all dated pieces.

Considering the very extensive amount of compositions that Telemann's *œuvre* comprises, these examples give an impression of the time-consuming extent of the corpus analysis work. Up to now it is nearly impossible to cover complete corpora (especially if they comprise more than one composer). Musicology is forced to work in a random sampling manner.

Accordingly, it is laborious and difficult to formulate and investigate new hypotheses. The expenditure of time prevents a systematic exemplification, verification and quantization of hypotheses on more comprehensive corpora or the complete amount of works of a composer.

Given an adequate machine-readable symbolic representation (e.g., MIDI, MusicXML), computer-driven analysis and data mining techniques could be able to include the big amount of data/big corpora in an analysis-run. In this way it can also be possible to discover new stylistic features and relationships between individuals within the analysis set.

These techniques of musical style analysis have to incorporate harmonic, rhythmic, motivic/thematic, key signature/tonal and contrapuntal features. These include aspects of modulation, motif variation, the treatment of dissonances, as well as phenomena of the voice-leading and consequential relations between parts.

Furthermore, the sung text/libretto and its musical treatment give enlightening clues for stylistic assignments. Here, the composers' musical vocabulary correlates to non-musical content and reveals his very specific personal way of compositional treatment.

Most of these aspects require more than a formal or statistical analysis. Ambiguity (e.g., during a harmonic modulation), motivic relationship and singular statistically insignificant phenomena are usually beyond the computational. But since these are the compositional means to work out the inner-musical structure—to accent and to catch the attention of the audience—these are usually the most important aspects for music analysis. It requires the human self-conception of musical expressivity and the interpretation of meaning. Thus, an automatic music analysis cannot and should not replace the musicologists' manual analysis.

It should be a tool for it, guided by human interaction that corrects and refines the information retrieval and visualizes its results for further investigation and interpretation.

Stylistic analysis in computer science

Many applications in the new research area of Music Information Retrieval (MIR) rely on automatic feature extraction techniques either on audio data or symbolic representations.




Figure 1: Narmour patterns

The content of audio files or scores can be characterized by high-dimensional feature vectors. One fundamental problem in all these approaches is (i) to determine common characteristics of specific groups in a given corpus (classification) or (ii) to evaluate the similarity between elements and to formalize these findings in metrics. Based on classifiers, new elements can be assigned to an appropriate class (e.g., genre classification, automatic play list generation). Similarity metrics can be used to determine clusters of similar elements. Melody or rhythm search applications like *query by humming* (e.g., Ghias et al., 1995) and *query by beat boxing* (e.g., Kapur et al., 2004) require flexible similarity measures to find an appropriate element in musical databases for a small melodic fragment that highly differs from the stored original.

The research in music psychology has shown that the recognition of melodies is based on their interval structure and their contour (Dowling, 1978). Therefore, the majority of melody retrieval systems uses strong abstractions of the melodic properties: direction of movements in the melody (Ghias et al., 1995), the rhythm as absolute duration (McNab et al., 1996) or relative to the meter (Chai & Vercoe, 2002).

Melodies have often been used as examples in the Gestalt theory. Meyer (1956) suggested closures to explain organization principles in melodies. Meyer's theory inspired Narmour's Implication-Realization model (Narmour, 1990; Narmour 1992) which is based on closures—expectation of melodic continuations. According to the Gestalt Law of Good Continuation (Wertheimer, 1923) the recognition of two similar or distinct elements raises the expectation of a continued similarity or variation. These elements can be pitches, intervals or directions. Figure 1 presents some typical melodic patterns: the continuation of pitches (*duplicate*), intervals with a change in direction (*intervallic duplicate*) or directions (*process*). The pattern *intervallic process* is based on the observation that nearly the same starting pitch is reached at the end of the pattern. Finally, the pattern *reversal* can be considered as a closure based on diversity,

but it is also a typical strategy in voice leading to absorb a leap by a step in the opposite direction. Three other patterns can be created through an octavation of the final note.

Several melody search systems use Narmour's patterns to describe voice leadings (Grachten et al., 2005; Grachten & Arcos 2004). Recently, Narmour's theory has been extended with mathematical models of melodic tension (Margulis 2005; Farbood, 2006). Figure 2 presents several features for the initial motiv of Franz Schubert's "Sängers Morgenlied".

Similarity measures for melodies are frequently based on the edit distance (Levenshtein, 1966), a metrics that determines the minimal number of operations that are required to transform one representation into another (insertion, deletion, substitution).

Fortunately, several powerful toolkits are available to extract features from audio data and symbolic representations. Very popular is the Humdrum toolkit for the ****kern** format (Huron, 2002) or jSymbolic for MIDI (McKay & Fujinaga, 2006). Moreover, the WEKA data mining toolkit (Witten & Frank 1999) offers many classification algorithms (e.g., decision trees) and clustering of groups in the feature space (e.g., k-means, k-nearest-neighbor, decision trees and support vector machines).

Computer science in musicology

One interesting direction in interdisciplinary musicology is the application of machine




Figure 2: Melody analysis of Franz Schubert's "Sängers Morgenlied"

learning techniques to the task of composer or region attribution. The majority of these algorithms was successfully applied to classify audio files accord to the genres (e.g., Dannenberg et al., 1997; Tzanetakis, 2001). Several researchers suggest to consider the scores of an individual composer as a class. Classification algorithms can then be applied in order to determine the characteristic features that distinguish the classes. Kranenburg & Backer (2004) used the C4.5 algorithm to construct decision trees automatically from training data. Their corpus contained scores from five composers of the Baroque and Classical era that were labeled with the name of the composer. In this corpus the feature vectors of the different compositions of the composers tend to form clusters in the feature space. Therefore, clustering algorithms (e.g., k-means) were able to detect these clusters automatically. Moreover, the decision tree provides insights into the personality styles of the composers. Jürgensen & Knopke (2004) applied classification techniques to determine the author attribution and origin in a corpus with many anonymous compositions.

Despite of the multitude of approaches to design flexible and robust algorithms to search for melodies, almost no papers focus on the detection of motifs. Apparently, the variety of variations—motivic developments—raises several problems for an automatic detection. The only approach to detect motifs was proposed by Lartillot (2007).

All these approaches to determine style features can benefit from visualizations that enable experts to interactively search for interesting phenomena (e.g., correlations between features, clusters, outliers) in large datasets. Both Kranenburg & Backer and Jürgensen & Knopke use scatter plots to visualize dependencies between style features. There is still a lack for expressive visualizations for musical data mining. Self-organizing maps or Kohonen networks (Kohonen, 1995) are powerful mechanisms to detect clusters in feature space and to map high-dimensional feature spaces to two dimensions (see Toiviainen & Eerola, 2006). But research in Information Visualization (e.g., Card et al., 1999, Ware, 2000) or Visual Analytics has shown the potential of an interactive exploration of high-dimensional data. Experts should be able to inspect different features in separate views, and select interesting data segments while these sections are also applied on all views (brushing). But musical structures also raise new challenges for visualizations. Sapp

(2005), for instance, presented a new visualization of modulations within longer scores.

Despite the fact that many research problems in musicology require extensive corpus analyses, the potential of data mining techniques to support the validation of musicological hypotheses or to detect the characteristics of individual scores or groups of compositions is currently not known in musicology. McKay's & Fujinaga's (2007) overview article is a good starting point to overcome this lack.

Data mining and visualization in musicology

Dataset

The Telemann Research Center hosts a large archive of partly unpublished and undated scores of Georg Philipp Telemann. In a long-time effort, this corpus has been partly transformed into a symbolic representation (DARMS format, Dydo, 1997). These scores were manually created and corrected in a period of more than 10 years. Due to the restricted printing capabilities in the Note-Processor (e.g., figured bass, articulation and dynamics), the final print-outs are enhanced with this information.

In order to achieve high-quality results for publications and in order to ease the editing process, a conversion to a more flexible musical notation format was necessary. We have chosen MusicXML (Good, 2001), as this format is well documented and all major music notation programs offer import and export filters for this format. Our music analysis program can read scores in MusicXML, CapXML and MIDI. As enharmonically equivalents cannot be distinguished in the MIDI format, our import filter determines the underlying scale of the score. This broad variety of import formats enables us also to apply our data mining tools to other corpora or to use other databases of sheet music (e.g., Mutopia¹).

Iterative data mining and visualization

Figure 3 illustrates our perspective of applying data mining techniques in musicological corpus analyses. An expert selects scores and applies automatic *feature extraction* techniques in order to evaluate some hypotheses or to detect interesting patterns in the corpus. Subsequently, the result for individual scores or groups of scores can be inspected.

Therefore, *visualizations* of the analysis have to be provided. For the visualization of some

aspects of analysis the results are directly integrated into the score. The features are assigned with their corresponding entities in the score. Subsequently, export filters transform this internal representation into the LilyPond music notation format.²

Finally, experts can correct errors, resolve ambiguities or insert new interpretations (*data enhancement*). Thus, contextual factors can be considered. Experts should also be able to insert annotations directly into the score. In our vision, a musicological analysis involves repeatedly iterated applications of data mining techniques. Moreover, a corpus may even contain several (conflicting) interpretations that encode subjective points of view.

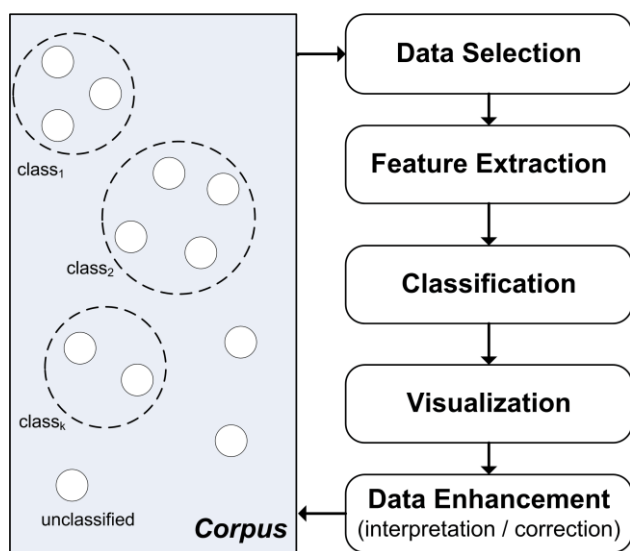


Figure 3: Application of data mining techniques in a musical corpus

Musical features

Many authors distinguish between low-level features that are derived from physical properties of audio files, and high-level features that are associated with properties derived from a symbolic representation. Our system only extracts features from symbolic representations. Therefore, with low-level features, we refer to properties that are directly associated to elements of the score. The detection of these high-level correlations requires additional inferences and contextual information.

Musical features are associated to the melodic, rhythmic, harmonic and lyrical properties listed below:

Melody: pitches, intervals, melody contour, step vs. leap, Narmour pattern (see Figure 1), ambitus/range of individual parts, pitch and interval histogram, di- and trigrams over all features;

Rhythm: absolute duration of notes and rests, relative duration in beats (with respect to the time signature), duration histograms and di- and trigrams over all features;

Harmony: chord, chord inversion, base note, highest note, duplication of pitches in the chord, chord histogram, chord di- and trigrams;

Lyrics: text.

The harmonic analysis is based on a chord dictionary that is constructed on the basis of the available scales and the interval structures of chords. This algorithm implements the chord scale theory (see Nettles & Graf, 2006). In this way, all chord variants and inversions are constructed. Our current system comprises

- 67 scales-based tone material with five, six, seven and eight tones and their modes and
- the interval structures of triads and sixth, seventh and extended chords. (major 4th, ninth, eleventh, thirteenth).

Hence, all triads are constructed from an interval pattern [1 3 5]. An additional dictionary specifies several descriptors for chords according to their interval structure. The chord pattern [1 3b 5b], for example, is constructed from a harmonic minor scale and is referred as *diminished* or *:dim* in LilyPond's terminology. This systematic way of chord construction also guarantees that chord synonyms (e.g., C-major-7 and A-minor-6) are detected automatically. Within the harmonic analysis of a score, the interval structure of the chords is compared with both chord dictionaries in order to determine chord descriptors.

Even the analysis of low-level properties results in a high number of features. McKey & Fujinaga (2006), for example, list 160 features in their MIDI-based feature extractor.

In our project, we are especially interested in high-level features (e.g., characteristics of voice-leading and the treatment of harmony

progression). Therefore, we also determine parallels in fifths and octaves between parts, the distances between voices, and the treatment of dissonances (anticipations, neighbor tones, passing tones, suspensions, escape tones and pedal points). The analysis of non-chord tones is based on the rhythmic structure of the score and adds new chord hypotheses to the harmonic analysis. Finally, all chord hypotheses are ranked according to the scale analysis.

The detection of several musical features (scale, mode, key signature, modulation) is very often difficult and can benefit from manual specifications. Therefore, experts can provide hints. These manual specifications are considered within subsequent feature extraction algorithms. This procedure enables experts to share their expertise and control all parameters of the data mining process.

Interactive feature search

The discussions in our project between experts in history of music, visualization and data mining revealed the need to support interactive queries that may combine aspects of melody, rhythm, harmony and lyrics. One example is the search for an individual motif and its development in the corpus. In the current prototype, experts can specify melodic and rhythmic patterns; in response to these motif queries, the system determines all variations of motifs (e.g., transposition, reversal of intervals and direction, note splitting, augmentation, diminution). With the following query, e.g., one can search for the initial motif in Franz Schubert's "Sängers Morgenlied" (see also Fig. 2):

```
c4 g8 | e4 c8 f4. g16 a8 f-sharp4 g8
```

Note that this motif specification also comprises a prelude. The system determines all musical segments that are somehow similar. Therefore, we apply the similarity measures based on the pitch contour, Narmour patterns, and motif/thematic structures.

Implications

Today, music analysis is still an extensive manual work. A highly adaptive computer-aided music analysis offers possibilities to support the teaching and learning of compositional techniques; to evaluate the novelty or accordancy of musical pieces with respect to a manually selected corpus. Our data mining techniques are not restricted to Western or tonal music. Hence, contrastive analyses

within large corpora may enable us to study more complex musicological problems or can make expensive manual analyses more efficient.

The interactive analysis of musical artifacts from a musicological perspective raises new challenges for an adaptable and iterative analysis of large corpora within a contrastive analysis. In this scenario, data mining, machine learning and visualization techniques enable researchers to evaluate hypotheses from a broad range of perspectives (e.g., music history, music theory, ethno-musicology, etc.).

Future Work

Further research and development will focus on two directions:

Automatic analysis

The extension of our catalog of compositional means for structuring and expression can allow more complex analysis queries. Our aim is to widen the stylistic range of analyses to enable the application on very different and inhomogeneous corpora. This could open up the way for contrastive analyses of composers and style epochs and raises the question of how to measure stylistic difference or similarity.

Interface and output

Style descriptive formalisms are needed; these have to represent more than just contrapuntal rules or statistically significant frequency. Lots of automatic music composition systems have shown the tendency to generalization conflicts with the intention to statistically less significant features that make a piece of music interesting and worth to be analyzed (Hörnel & Menzel, 1999).

These features will be stored as meta-data and can be used for graphical output and as input by several further systems, for instance, to enable a listening analysis, a performance system can use this meta-information, add several means of expression according to them (dynamic, articulation, tempo, etc.), and realize an expressive performance.

Conclusion

We developed a flexible music retrieval system that experts in music history can use to search for melodic, rhythmic, harmonic and lyrical features that may confirm the objective of analysis. Color codes and textual annota-

tions are used to integrate the search results directly into automatically generated scores.

Telemann's oeuvre reflects the musical development from baroque to the classical period of European music. Thus, Telemann used certain musical features solely in specific periods of time. Therefore, our approach offers completely new perspectives for the dating of scores and a critical judgment of anonymous or controversial scores. Moreover, the musical treatment of the underlying lyrics (libretto) in vocal works has not been studied within the field of 'artificial intelligence and music', but raises interesting musicological questions.

References

- Card, S., Mackinlay, J., and Shneiderman, B., editors (1999). *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann.
- Chai, W. and Vercoe, B. (2002) Melody Retrieval on the Web. In *SPIE Multimedia Computing and Networking*, pp. 226–241.
- Dannenberg, R., Thom, B., and Watson, D. (1997). A Machine Learning Approach to Musical Style Recognition. In *Int. Computer Music Conf.*
- Dowling, W. (1978) Scale and Contour: Two Components of a Theory of Memory for Melodies. *Psychological Review*, 85(4):341–354.
- Dürr, A. (1952) Zur Echtheit einiger Bach zugeschriebener Kantaten. In *Bach-Jahrbuch 1951/52*, pp. 30–46.
- Dydo, S.J. (1997) DARMS: The Note-Processor Dialect. In Selfridge-Field, E. (ed.) *Beyond MIDI—The Handbook of Musical Codes*. MIT Press, Cambridge.
- Farbood, M. (2006) *A Quantitative, Parametric Model of Musical Tension*. PhD thesis, MIT Media Laboratory.
- Fleischhauer, G. (1963) Einige Gedanken zur Harmonik Georg Philipp Telemanns. In *Beiträge zu einem neuen Telemann-Bild, Konferenzbericht der 1. Magdeburger Telemann-Festtage*, pp. 50–63.
- Ghias, A., Logan, J., Chamberlin, D., and Smith, B. (1995) Query by Humming — Musical Information Retrieval in an Audio Database. In *ACM Multimedia*, pp. 231–236.
- Grachten, M. and Arcos, J. (2004) Using the Implication/Realization Model for Measuring Melodic Similarity. In *16th European Conf. on Artificial Intelligence (ECAI)*.
- Grachten, M., Arcos, J., and Lopez de Mantaras, R. (2005) Melody Retrieval using the Implication/Realization Models. In *6th Int. Conf. on Music Information Retrieval (ISMIR)*.
- Good, M. (2001) MusicXML for Notation and Analysis. In Hewlett, W.B. and Selfridge-Field, E. (eds.) *The Virtual Score: Representation, Retrieval, Restoration*. MIT Press, Cambridge, pp. 113–124.
- Hirschmann, W. (1992) Die gewisse Schreibart. Gedanken zum Problem des Personalstils bei Georg Philipp Telemann. *Concerto*, 9:17–37.
- Hobohm, W. (1983) Vom lombardischen Rhythmus bei Telemann. In *Die Bedeutung Georg Philipp Telemanns für die Entwicklung der europäischen Musikkultur im 18. Jahrhundert*, S. 4–22.
- Hörnel, D. and Menzel, W. (1999) Learning Musical Structure and Style with Neural Networks. *Computer Music Journal*, 22(4):44–62.
- Huron, D. (2002) Music Information Processing using the Humdrum Toolkit: Concepts, Examples, and Lessons. *Computer Music Journal*, 26:15–30.
- Jürgensen, F. and Knopke, I. (2004) A Comparison of Automated Methods for the Analysis of Style in Fifteenth-Century Song Intabulations. In *Conf. on Interdisciplinary Musicology (CIM)*.
- Kapur, A., Benning, M., and Tzanetakis, G. (2004) Query-by-Beat-Boxing: Music Information Retrieval for the DJ. In *Int. Conf. on Music Information Retrieval (ISMIR)*.
- Kranenburg, P. v. and Baker, E. (2004) Musical Style Recognition — A Quantitative Approach. In *Conf. on Interdisciplinary Musicology (CIM)*.
- Kohonen, T. (1995). *Self-Organizing Maps*. Springer Verlag, Berlin.
- Lange, C. (2004) "Der Abgrund muß erzittern" — Naturgewalten in der Musik Telemanns. In "Telemann — der musikalische Maler". Konferenz anlässlich der 17. Magdeburger Telemann-Festtage 2004. (in print).
- Lartillot, O. (2007) Motivic Pattern Mining in Symbolic Music: Managing Multidimensionality and Transformations in Music. In *31st Ann. Conf. of the German Classification Society*.

- Levenshtein, V.I. (1966) Binary Codes Capable of Correcting Deletions, Insertions, and Reversals. *Soviet Physics Doklady*, 10(8):707–710.
- Margulis, E. (2005) A Model for Melodic Expectation. *Music Perception*, 22(4):663–714.
- McKay, C. and Fujinaga, I. (2006) jSymbolic: A Feature Extractor for MIDI Files. In *Int. Computer Music Conf.*, pp. 302–305.
- McKay, C. and Fujinaga, I. (2007) Style-Independent Computer-Assisted Exploratory Analysis of Large Music Collections. *Journal of Interdisciplinary Music Studies*, 1(1):63–85.
- McNab, R., Smith, L., Witten, I., Henderson, C., and Cunningham, S. (1996) Toward the Digital Music Library: Tune Retrieval from Acoustic Input. In *ACM Digital Libraries*, pp. 11–18.
- Meyer, B. (1956) *Emotion and Meaning in Music*. Chicago University Press. Chicago.
- Narmour, E. (1990) *The Analysis and Cognition of Basic Melodic Structures*. University of Chicago Press, Chicago.
- Narmour, E. (1992) *The Analysis and Cognition of Melodic Complexity*. University of Chicago Press, Chicago.
- Nettles, B. and Graf, R. (2006) *Chord Scale Theory and Jazz Harmony*. Advance Music, Rottenburg.
- Reipsch, B. (1998a) Zum "Jahrgangsdanken" Georg Philipp Telemanns — dargelegt an Beispielen geistlicher Kantaten Eisenacher Dichter. In "Nun bringt ein polnisch Lied die ganze Welt zum Springen". Telemann und Andere in der Musiklandschaft Sachsens und Polens, (Arolser Beiträge zur Musikforschung, Band 6), pp. 63–76.
- Reipsch, R.-J. (1998b) Telemanns Rezitativtechnik — Untersuchungen am vokalen Spätwerk. In Bericht über den 11. Internationalen Kongreß der Gesellschaft für Musikforschung "Musik als Text" (Bd. 2: Freie Referate), pp. 283–294.
- Sapp, C. (2005) Visual Hierarchical Key Analysis. *Computers in Entertainment*, 4(4).
- Telemann, G.P. (1767) Letzte Beschäftigung Georg Philipp Telemanns, im 86sten Lebensjahre, bestehend in einer musikalischen Klang-und Intervallen-Tafel. Hamburg. (reprint in: Rackwitz, W. (ed.) *Georg Philipp Telemann. Singen ist das Fundament zur Music in allen Dingen. Eine Dokumentensammlung*, S. 266–273, 1981).
- Toiviainen, P. and Eerola, T. (2006) Visualization in Comparative Music Research. In *Computational Statistics (COMPSTAT)*, pp. 209–221.
- Tzanetakis, G. (2001) Automatic Musical Genre Classification of Audio Signals. In *Int. Conf. on Music Information Retrieval (ISMIR)*.
- Ware, C. (2000). *Information Visualization: Perception for Design*. Morgan Kaufmann.
- Wertheimer, M. (1923) Untersuchungen zur Lehre von der Gestalt II. *Psychologische Forschung*, 4:301–350. (also in Ellis, W. (1938) *A Source Book of Gestalt Psychology*. London: Routledge & Kegan Paul, pp. 71–88)
- Witten, I. H. and Frank, E. (1999) *Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations*. Morgan Kaufmann.

¹ *The Mutoxia Project: Free Sheet Music for Everyone.*
<http://www.mutopiaproject.org/>

² LilyPond ... music notation for everyone.
<http://lilypond.org/>