journal of interdisciplinary music studies
spring 2010, volume 4, issue 1, art. #10040102, pp. 17-43

Collaboration Perspectives for Folk Song Research and Music Information Retrieval: The Indispensable Role of Computational Musicology

Peter van Kranenburg[1], Jörg Garbers[1], Anja Volk[1], Frans Wiering[1], Louis P. Grijp[2], and Remco C. Veltkamp[1]

[1] Utrecht University
[2] Meertens Institute, Amsterdam

Background in Folk Song Research (FSR). In the past century melodic variation caused by oral transmission has been studied within the discipline of Folk Song Research (FSR). Also, various systems have been developed to categorize large collections of folk songs. Since the 1940s many attempts have been made to design automatic systems to categorize melodies. However, after several decades, no strong theories of oral transmission, and no generally applicable classification systems have yet emerged. Currently, many cultural heritage institutions give high priority to the digitization and unlocking of their (musical) collections.
Background in Music Information Retrieval (MIR) and Computational Musicology (CM). In the field of Music Information Retrieval (MIR) methods are designed to provide access to large bodies of music, while in the field of Computational Musicology (CM) computational methods are designed to study musicological questions. These developments stimulate a new interest in the questions of FSR. However, very few CM and MIR studies take the particularities of orally transmitted melodies into account.
Aims. Better collaboration between MIR, CM and FSR (or Musicology in general) will enrich Musicology with new methods to study existing problems and MIR with better understanding of music.
Main contribution. By surveying relevant achievements of the disciplines, we show a gap between MIR and FSR. To bridge that gap we provide promising directions for research based on current developments, as well as a collaboration model in which CM serves as an intermediate between FSR and MIR.
Implications. MIR should go beyond provided 'ground truth' data in implementing and testing models that generated those ground truths. The concepts used in folk song research, 'tune family' in particular, should be modelled, providing MIR a musically informed implementable model and FSR an enriched understanding of those concepts.

Keywords: Computational Musicology, Music Information Retrieval, Folk Song Research, Tune Family, Tune Classification, Tune Identification.

• *Correspondence:* Peter van Kranenburg; Dept. of Information and Computing Sciences; Utrecht University; P.O. Box 80.089; NL-3508 TB Utrecht, Netherlands; e-mail: petervk@cs.uu.nl
• *Action Editor:* Rytis Ambrazevičius
• *Received:* 26 February 2009; *Revised:* 10 September 2009; *Accepted:* 05 October 2009
• *Available online:* 30 December 2009
• doi: 10.4407/jims.2009.11.006

1. Introduction

Recent developments in Musicology include a growing interest in empirical approaches aiming to enrich traditional qualitative research with data-rich quantitative studies (see e.g., Clarke and Cook 2004 or Huron 1999b). This approach can shed new light on the objects of interest of Musicology, such as musical artifacts, their interrelations and their relations to human culture and behavior. This quantitative research is facilitated by recent improvements in computer technology enabling the use of computationally intensive methods. The current article explores the promises and problems of a computational approach to the study of folk song melodies.

Folk songs are sung by common people during work or social activities. One of the most important characteristics of these songs is that they are part of oral culture.[1] The melodies and the texts are learned by imitation and participation rather than from written sources such as books. In the course of this oral transmission, changes occur to the melodies, resulting in groups (‘tune families’) of more or less related melodies.

During the second and third quarter of the twentieth century research on this kind of music flourished in the field of Folk Song Research (FSR). Many folk songs were recorded on tape or transcribed on paper and are thus available for research. Various attempts were made to find structures and patterns in the various folk song corpora. However, after several decades, no strong theories of oral transmission, and no generally applicable classification systems have emerged.

Recent developments in Computational Musicology (CM) and Music Information Retrieval (MIR) have the potential to facilitate and support the research on folk song melodies. CM studies musicological questions with computational methods, while in the field of MIR tools are being developed to unlock large bodies of music. Providing new research methods, these developments stimulate a new interest in the questions of FSR.

There are at least two important reasons for employing MIR technology in FSR. Firstly, the musical models that have to be developed in MIR to process large amounts of folk song data are relevant to the study of the folk song melodies themselves and their histories. Secondly, MIR technology is of invaluable importance for the preservation and unlocking of large melody collections. (Ethno)musicological archives contain the musical ‘memory’ of the world. Therefore means for maintaining and accessing these archives are necessary. Currently, many cultural heritage institutions are giving high priority to the digitization and unlocking of their collections, including musical archives.[2] Hence, the development of computational means to do so is an urgent matter.

Although attention has been paid to folk songs in the MIR community, very few studies focus on the particularities of orally transmitted melodies. In most cases folk songs were simply used because they were available as a test collection. Serious attempts to build software for processing folk song melodies should model concepts and methods that were developed in FSR. However, this is not yet standard practice. Major impediments for fruitful collaboration are the unfamiliarity of researchers in both fields with each other's methods and traditions, and the non-formalized nature of many FSR concepts and theories. Therefore we need to find an approach to bridge this gap.

This article provides overviews of approaches to the study of folk song melodies in Folk Song Research, Computational Musicology and Music Information Retrieval (sections 2 and 3), identifies the problems that arise when methods from Computer Science are applied to research questions of FSR (section 4), presents concrete directions for research (section 5) and describes a collaboration role model (section 6). Although the current focus of interest is on folk songs, most of the questions that have to be resolved to conduct a successful computational research project to folk songs do also play a role in wider scopes. Therefore, this article can be read as a case study in practicing and reflecting on Computational Musicology.

This article was written in the specific context of the WITCHCRAFT project (Wiering 2009).[3] The goal of the WITCHCRAFT project is to develop a content-based retrieval system for folk song melodies stored as audio and notation. This system will give access to the collection of folk song recordings and transcriptions of the Meertens Institute (a research institute for Dutch language and culture in Amsterdam). It will enable content based retrieval for the Database of Dutch Songs, which already has a fully functional meta data search engine.[4] Its purposes are on the one hand to support Folk Song Research (FSR) in classifying and identifying variants of folk songs and on the other hand to allow the general public to search for melodic content in the database of the Meertens Institute. In the current paper we focus on the former purpose.

2. Some Past and Current Approaches to Melody in Folk Song Research

Since the late nineteenth century, the availability of collected folk song melodies has generated a considerable amount of musicological research. One of the primary concerns is how to deal with the specific type of melodic variation caused by the process of oral transmission. The basic question is how to model the relationships between melodies from the same folk song culture. Therefore we will first characterize oral transmission, and then classification and identification of melodies in the context of FSR will be briefly discussed.

2.1. Oral Transmission and Tune Families

The capabilities of such human faculties as perception, memory, performance and creativity play an important role in the transmission of songs in an oral tradition. Performers have more or less abstract representations of songs in their memories. The only way in which others have access to a song is to listen to a performance. Research into music cognition (Peretz and Zatore 2005) shows that the representation of a song in human memory is not 'literal'. During performance the actual appearance of the song is reconstructed or recreated. In the process of transforming the memory representation into audible words and melody, considerable variation may occur. As long as the processes of encoding songs in, and performing songs from human memory are not sufficiently understood, we have to focus mainly on the recorded or transcribed song instances in order to infer knowledge about this kind of variation.

This approach was taken by Walter Wiora (1941), resulting in a comprehensive inventory of types of variation in German folk songs. Wiora summarizes the issue as follows: "Alles an der Beschaffenheit einer Melodie ist veränderlich".[5] He divides the types of changes into seven categories: 1. changes in contour, 2. changes in tonality, 3. changes in rhythm, 4. insertion and deletion of parts, 5. changes in form, 6. changes in expression, and 7. demolition of the melody. He provides many examples for each of these types of change.

The difficulty of understanding the kind of melodic variation caused by oral transmission is clearly stated by Bertrand Bronson (1951, p.51):

> "When we consider that there is no accessible original to impose its authority; that at every moment in its history such a tune is open to all the gusts of casual influence, subject to forgetful recollection, to individual, or local, or epochal, preferences in mode and rhythm, to willful invention or derangement, to the accidents of marriage with continually altered verbal patterns that impose their own necessities upon melodic statement, and all these operative without any counterbalancing overt external control; we can only marvel at the inner urgency with which folk tunes maintain their essential selfhood in the face of such overwhelming odds."

Together with the remark of Wiora that everything in a melody can change, this citation suggests that the melodic variation caused by oral transmission is a holistic process. Therefore, it cannot be understood by only considering a selective number of melodic features. This is confirmed by another remark of Bronson (1949, p.169):

> "All who have worked with the problems of variation in a related body of [melodic] materials will readily acknowledge that the question of relatedness involves far more than a mere note-for-note or accent-by-accent correspondence. One very soon comes to realize that this is a problem of the utmost subtlety, in which potentially are included all or most of the elements constituting melodic identity; range, melodic and rhythmic mode, number of phrases, patterns of phrasal combinations, refrain schemes, cadence points, and so on to minuter particulars. It therefore becomes desirable to establish, at least tentatively, the relative weight to be allotted to some of these elements. It may well be that herein, with due discrimination, we shall ultimately find the basis of just distinctions between "families" and more inclusive patterns, or "types" of melody."

For a corpus of British-American folk songs Bronson (1950) obtained weights for some aspects of melodic identity. But he did not test this ordering of importance on corpora from other cultures.

We will now present some attempts from Folk Song Research to better understand the process of melodic change that is caused by oral transmission.

Ernst Klusen *et al.* (1978) conducted an experiment in which melodies were passed orally from one person to another. Their most general finding is in accordance with the result of Wiora: every tone in a melody can change, but some tones are more stable than others. In general, pitch was more varied than rhythm. However, as they indicate themselves, their experiment only tested relatively short-term memory (a few weeks).

The concept of *tune family* was developed by Samuel Bayard (1950) and defined as: "a group of melodies showing basic interrelation by means of constant melodic correspondence, and presumably owing their mutual likeness to descent from a single air that has assumed multiple forms through processes of variation, imitation, and assimilation." Bayard supposed that the entire body of Anglo-American folk songs consists of forty or fifty such families (Nettl 2005, p. 116).

After studying traditional music of Ireland, James Cowdery (1984) proposed a "fresh" view on the concept of tune family. He criticizes Bayard's tune family model by posing that folk musicians do not compose new melodies as new instantiations of abstract archetypical airs, but relate new melodies to other concrete melodic material they know. Therefore, Cowdery does not only focus on global similarity but also on motifs that are shared among members of a tune family. Melodies may have some sections in common while other sections differ. He proposes three principles for studying relationship between melodies of which the first one corresponds to Bayard's definition: 1. the "outlining" principle: melodies correspond in their overall contour, 2. the "conjoining" principle: melodies have sections in common, while other sections differ, and 3. the "recombining" principle: melodies are composed from material from the same "pool" of melodic motifs. Thus there is no absolute need for the hypothetical historical sequence of melodies, which in virtually all cases has been lost, if existed at all. Instead, all the melodic material used to analyze melodies and to relate them by means of the tune family concept is concrete melodic material that can be observed and that is meaningful to the folk musicians.

The approaches presented so far are based on the hypothesis that a historic link between two melodies implies some kind of melodic similarity. To show that this is not necessarily the case, Bruno Nettl (2005, p.116) describes a way in which a song may change entirely. Starting with the structure ABCD, the first half of the song may be 'dropped', leaving CD, which may be extended with new material resulting in CDEF, which may end up in EFEF by dropping the first part again. In this case there is a historic, but not a melodic link between the first and the last song.

There is no generally accepted theory that explains melodic variation in oral cultures yet. David Rubin (1995) has developed a cognitive theory for oral transmission of texts in which variation is modelled by constraint-based reconstruction of the texts from memory. For example, two words with the same meaning and the same metrical characteristics may both be found at corresponding places in a set of variants of a text. The actual appearances of variants may change, but they do obey to the same constraints to a certain extent. A similar process might occur during transmission of melodies. To exploit this for retrieval purposes, the challenge is to model these constraints.

2.2. Identification

If two song instances are derived from the same common 'ancestor', they belong to the same tune family (Bayard 1950) and are considered to be the same song, or, more precisely, manifestations of the same song (Nettl 2005, p. 114).[6] The identity of a song is a complex and abstract concept. It is not obvious what constitutes the 'substance' – or, in the words of Bronson, the essential selfhood – of a song that is shared among historically derived variants. As a consequence, in folk song classification systems that are based on a limited number of features, historically linked variants may erroneously end up in entirely different classes. The possibility of interference between tune families complicates the issue even further. Because the concept of identity goes beyond individual features of song instances, it is very difficult to develop models that explain tune families.

However, identification of melodies is necessary to address a number of research questions, such as: Where do the individual songs originate from? What were the most popular melodies in a certain time or at a certain place? Which influences from abroad can be traced? How did the melodies develop over time?

Because folk song collections contain only a sample of the melodies and variants that have existed, it is impossible to find all variants that are derived from a common 'ancestor' melody and thus to reconstruct the complete history of melodies from the material in the collection. However, in many cases it is feasible to find related groups of melodies within the collection, based on melodic and textual similarity and available metadata, and to try to link them to tune families as a second step. For this a retrieval system can be an important tool.

Identification of melodies is also important for improving access to cultural heritage. For publications like folk song books, CD-boxes, etc., it is important to know which melodies are in a collection and what their relations and identities are. It is unfeasible to find all relations by hand; therefore a computational approach is desirable.

2.3. Features for Classification

In FSR, classification systems are used to put melodies in some kind of rational order. These systems are based on such features as the number of lines, the number of syllables or the cadence note sequence of a song.[7]

The purpose of the most important classification systems in FSR has been twofold. Firstly, classification is desirable for e.g. storing melodies in a card file database or for publishing a book with melodies. In those cases a one-dimensional ordering is required. Such an ordering must provide an easy way to retrieve a melody. Currently, this necessity has been overcome by using digital databases, but this requirement has been important for classification methods that are still in use. Secondly, one of the main aims of these classification systems is to group melodies that are related through the process of oral transmission together (Nettl 2005, p. 123). Hence, a classification system can be considered a hypothesis of how melodies relate to each other in the process of oral transmission, or as a practical tool to identify melodies. However, in FSR formalized tests of classification systems with respect to their ability to group melodies from the same tune family have not been performed.

Here we give a selective overview of features used in various folk song classification systems. More complete surveys can be found in Elscheková (1966), Keller (1984) and Bohlman (1988).

Figure 1. The Dutch song *Dat gaat naar Den Bosch toe* as notated in Brandts Buys (1975, p.81). For the first phrase, accented notes on various levels are marked.

Most classification systems were developed for specific corpora. One of the first was developed for Finnish songs by Ilmari Krohn (1903). In his system the number of lines and the cadence notes (ending notes of the lines, as depicted in Figure 1) are most important. Béla Bartók and Zoltán Kodály adapted his system for Hungarian folk songs. In their publications songs were hierarchically ordered by: 1. the number of lines, 2. the sequence of cadence notes, 3. the number of syllables in each line, and 4. the range (Suchoff 1981, p. xxxiv). In later work, Bartók used another system in which he divided Hungarian songs into three classes, namely old style, new style and mixed style melodies (Suchoff 1981, p. xlii). Subdivisions were made according to rhythmic characteristics and the number of lines. Obviously, this way of ordering the material is specifically aimed at the corpus of Hungarian songs. As Bruno Nettl (2005, p. 124) points out, Bartók's particular choice of features for classification could only be made by someone already familiar with the corpus for which the system was developed. This applies to folk song classification systems in general (Bohlman 1988, p. 33).

For the British-American folk song tradition, Bertrand Bronson (1950) ordered a list of features according to importance using a punch card system: 1. final cadence, 2. mid cadence, 3. first accented note, 4. first phrase cadence, 5. first accented note of second phrase, 6. penultimate stress of second phrase, etc. Thus, a classification system based upon these features could be expected to group songs in the same tune family together to some extent.

In an important publication of The German Archive of Folk Song (Deutsches Volksliedarchiv) containing German ballads (Suppan and Stief 1976), an ordering is used that reflects the system of Krohn. The first criterion is the number of lines and the second criterion is the cadence note sequence. If the editors judged a resulting group incoherent, subgroups were made. The exact way in which these subgroups were established, is not accounted for.

3. Computational Approaches to Folk Song Melodies

In this section we survey research results from Computational Musicology and Music Information Retrieval concerning the study of folk song melodies, as well as melody search engines that have large numbers of folk song melodies in their databases.

3.1. Folk Songs in Computational Musicology

For the overview in this section the field of Computational Musicology is taken in its broadest sense: any research on melodies that makes use of computational methods in one way or another.

3.1.1. The Early Days

Ordering melodies according to some specific criteria is to a certain extent a ‘mechanical’ activity. Therefore it is not surprising that the use of computer systems was considered soon after they became available. As early as in 1949, Bertrand H. Bronson proposed a method to represent folk songs on punch cards. Thus songs with certain desired characteristics could be retrieved using a sorting machine (Bronson 1959). In the following decades many studies on the use of computers in folklore and folk music were published. A bibliography from 1979 on this topic lists more than 100 references (Stein 1979).

Starting in the 1980s, an enormous digitization project was carried out under supervision of Helmuth Schaffrath. He developed The Essen Associative Code (EsAC), a music encoding for monophonic folk songs. A detailed description is provided in Schaffrath 1997. After one and a half decade, more than 14.000 songs were digitized. The collection is still being used as test collection for melody analysis or melody retrieval. Along with the collection analytical software was implemented, which could extract numerous features from the songs, such as distributions of intervals and durations, rhythmic patterns, cadence tone sequences, pitch contour, etc. Schaffrath (1992) shows how to use these features to retrieve melodies from a database of folk song melodies.

In the 1980s Wolfram Steinbeck (1982) and Barbara Jesser (1991) did research into computer aided analysis of monophonic music. Both used a subset of the Essen folk song collection. Jesser evaluated statistics of features such as interval frequencies, duration frequencies, range, accent and cadence tone sequences and other features. In a number of example tune families, she showed that in each family common characteristics could be found, but that it was not the case that for all families the same features are important.

Wolfram Steinbeck (1982) focused primary on clustering. He used hierarchical clustering algorithms to group melodies from the Essen collection together using 13 features such as mean and standard deviation of pitches, range, size of intervals, number of direction changes, and others (Steinbeck 1982, p. 275). With a set of 35 melodies he was able to cluster melodies into meaningful groups, such as hymns, children's songs and hunting songs. An experiment with 500 melodies also led to clusters, but in this case the clusters were more difficult to characterize musically.

3.1.2. Contour-Based Approaches

There are some contour-based approaches. Using tools from his Humdrum Toolkit (Huron 1999a), David Huron (1995) confirmed the hypothesis that folk song melodies tend to show arch-like contours in single lines as well as in successions of lines. His testing material consisted of melodies from the Essen collection.

Zoltán Juhász (2000, 2002, 2004, 2006) published several studies using a large collection of digitized Hungarian folk song melodies. In his approach, a melody is represented as a contour vector, which is constructed by sampling the pitch at equal distances in time. If the sampling frequency is high enough, all details of the contour are preserved in the resulting vector. Juhász performed a principal component analysis of the contour vectors of melodic phrases.[8] The principal components can be interpreted as contours themselves. In the space spanned by the first few principal components, clusters of melodies can be found that share the same contour characteristics, namely, clusters of phrases with the same beginning and ending notes (2002). He also compared the Hungarian melodies with songs from other countries in Europe by training Self Organizing Maps with the contour vectors (2006). By evaluating the extent to which the map of one culture is excited by contour vectors from another culture, one can evaluate the extent to which the cultures share contour types. It appears that all involved musical cultures have some contour types in common. This caused Juhász to speculate about a "common language" that reflects an archaic common origin of these European traditions.

3.1.3. Segmentation

Some segmentation algorithms that have been developed within CM were aimed for or tested on folk song melodies. Zoltán Juhász (2004) used the entropy of the continuation of a sequence of intervals. A high value of this conditional entropy implies that the next interval is hard to predict, which may indicate a segment boundary. This approach can be used for segmentation.

Another data-driven approach to segmentation is presented by Rens Bod (2002). The segmentation is performed by a parser following rewrite rules that are inferred from a training set, in this case a part of the Essen collection. The rewrite rule with the highest probability is followed. He was able to reproduce 81% of the line breaks that were encoded in the Essen collection.

3.1.4. Other Approaches

Darrel Conklin and Christina Anagnostopoulou (2001) used the concept of viewpoint. This is a formalized way to represent a feature of a melody, by means of a sequence of symbols. Using a contour viewpoint and association rule mining Conklin (2006) could confirm the finding of Huron (1995) that melodies tend to have an arch shape. In combination with suffix trees, viewpoints can also be employed to find repeating patterns.

Bret Aarden and David Huron (2001) proposed the use of geographical information. Thus geographical variance of some features could be visualized by showing densities on a map.

The studies described show various valuable approaches to the processing of folk song melodies. With exception of those by Zoltán Juhász, none of these studies explicitly state an interest in folk songs as part of oral tradition. The particular questions of FSR, such as the understanding and modelling of tune families are barely addressed. Hence, we can state that the current available results from CS do not get us much further in addressing the problems of FSR.

3.2. Folk Songs in Music Information Retrieval

Folk song melodies have been used in a considerable number of MIR studies. However, in many cases folk songs were chosen because of their availability and not because of an interest in folk music as such. This applies to all 11 papers in the complete ISMIR proceedings from 2001-2008 that employ the Essen folk song collection.[9] In none of these papers the implications of the choice for this data set are discussed. In most cases it is just stated that the collection is used, or a pragmatic reason is provided, e.g., the need for a large music database, or the need for a collection of monophonic songs. The results of the more general questions addressed, such as metre classification, benchmark establishing or segmentation, have not been interpreted concerning their potential to contribute to folk song research.

3.2.1. Search Engines

Some online search engines allow the user to search in a large collection of folk song melodies.

The database of the **Danish Folklore Archives** contains about 10,000 instrumental melodies found in books and manuscripts with Danish folk music.[10] The collection can be queried on musical content with two kinds of accent note patterns: a so-called “incipit note sequence” (notes on the beats) and an “accent note sequence” (notes on the first beat of each bar), corresponding to the beat and the strong accents levels in Figure 1. String matching is used to evaluate the similarity of the query with the melodies in the database, with the number of permitted errors as parameter. In the explanation on the web site it is stated that the search for accent and incipit patterns in practice has proven to be the most reliable way to search a collection of melodic variants.[11] Apparently, these are supposed to be stable elements in Danish tunes by the creators of the search engine. However, a full account of the rationale behind the choice for accent and incipit note sequences is not provided.

The **Colonial Music Institute**, which promotes research in early American music and dance, offers an index for about 75,000 instrumental and vocal pieces from the period 1589–1839 (*sic*), including social dance tunes and songs.[12] From each melody an incipit is present in the database. There are three ways to browse these incipits: a sequence of scale degrees of all notes, a sequence of scale degrees of stressed notes, and a sequence of intervals.

Another online searchable database with folk song melodies is the **Digital Archive of Finnish Folk Tunes**.[13] The collection contains about 9,000 melodies, most of them collected in the early 20th century by Ilmari Krohn. Two types of melodic query can be used to search: gross contour (Parsons code: `"u"` (up), `"d"` (down), and `"r"` (repeat)) and a sequence of intervals in semitone distance. Wildcards may be used to allow a sub-pattern search. A string matching algorithm is used for matching.

There are also some search engines that are not specifically aimed at unlocking folk song collections, but nevertheless contain a large number of folk song melodies.

Themefinder can be used to search a large collection of themes and incipits from classical works, folk songs (the Essen collection; Schaffrath 1995) and 16th-century Latin motets.[14] Several representations of a query melody can be used: pitch sequence (e.g., `"G G G E-"`); interval sequence (e.g., `"P1 P1 -M3"`); scale degree (e.g., `"5 5 5 3"`); Gross Contour (e.g., `"ssd"`, in which `"s"` means 'same'); Refined Contour, in which steps and leaps are distinguished (e.g., `"ssD"`); Key and Metre. The user can choose whether the query should occur at the beginning, or anywhere in the theme. Internally the Humdrum Toolkit (Huron 1999a) is used to perform the search.

The melody index **MELDEX** has been developed in the context of The New Zealand Digital Library project at the University of Waikato (McNab *et al.* 1997).[15] The major part of folk songs in the Meldex database consists of the Essen folk song collection together with a collection of songs from the Digital Tradition Folk-Song Database.[16] Both text and melody queries are possible. A query melody can be generated by clicking keys on a virtual keyboard, or by whistling or humming in a microphone. An approximate string matching algorithm is used for matching (McNab *et al.* 1997). Unlike the previous discussed search engines, Meldex allows approximate matches without using wildcards or specifying the permitted number of errors. The user can choose whether the interval sequence or the contour of the query melody should be used. The inclusion of rhythm information in the search is optional and is switched off by default.

The database of the "Open Music Encyclopedia" **Musipedia** contains a large number of folk song melodies.[17] Various input methods can be used to create a query:

Lilypond code, contour (Parsons code), humming/whistling, or tapping. For matching, transportations distances (Typke *et al.* 2003) as well as string matching (for the Parsons code) are used. Among the described search engines, this is the only one in which rhythm is fully involved by default.

Figure 2. First lines of the search results. **Q** is the query. **R1** is found by the Danish engine, **R2** and **R3** by the Finnish engine, **R4** and **R5** both by Themefinder and Meldex, **R6** and **R7** by Musipedia, and **R8** by YahMuugle. All melodies are transposed to G major. The titles are: **Q** *Dat gaat naar Den Bosch toe*, **R1** air *Om al Verden er*, **R2** [without title], **R3** *Nelosta*, **R4** *Loot ons noch ens drinken*, **R5** *Ueber die Beschwerden dieses Lebens*, **R6** *Scottisch Simple de Guemene*, **R7** *I'm a little tea pot*, **R8** *Variations* by Aloys Schmitt. Two phrases from the original composition by Pierre Gaveaux, **G1** and **G2**, are added for comparison

Only the Danish search engine and the index of the Colonial Music Institute have query methods that are explicitly motivated by knowledge or hypotheses from FSR. One can search for a sequence of accented notes, which are assumed to be more stable across variants of a tune than unaccented notes.

3.2.2. An Example

As an example to show what can be accomplished by using the search engines described in the previous section, we will take a folk song melody as a query and discuss the search results.

In the Database of Dutch Songs,[18] the melody of the Dutch folk song *Dat gaat naar Den Bosch toe* as shown in Figure 1 has been identified as the song “Contre les chagrins de la vie” from the opera *Le petit matelot* by Pierre Gaveaux (1786). The first two phrases of Gaveaux’ song are shown as melodies G1 and G2 in Figure 2. As indicated in Figure 1, the Dutch melody consists of two sections, A and B. Thus, the structure of the song is ABABA. This melody is our query. The search results are summarized in Figure 2.

For the search engine of the Danish Folklore Archives, we have to transform the melody into an incipit or accent note sequence. The incipit note sequence is `"13516665"`. These numbers are the scale levels of the notes on the first eight beats. Querying the database with the permission of one error results in one hit, the air *Om al Verden er* as notated in a nineteenth century book with airs and dances that was property of Hans Jensen Hansen. This melody is labeled R1 in Figure 2. The B-section of this melody, which is not shown in the figure, is similar to the B-section of the Dutch song. As can be seen in the figure, this melody is encoded in 2/4 metre. Thus the accent note sequence of the 4/4 query, which is defined as the sequence of notes on the first beat of each bar, is not compatible with the accent note sequence of melody R1 in Figure 2. This complicates searching for accent note sequences. If we construct an accent note sequence of the query as if it were notated in 2/4 metre and permit one error, melody R1 is in the result list at rank three, after two unrelated hits.

The search engine of the Digital Archive of Finnish Folk Songs can be queried with an interval sequence representing each interval in half tone steps. For our query melody this would be `"+2+2+1+2+5-3"`. Searching for this sequence results in melodies R2, without title, and R3, entitled *Nelosta*, at ranks 1 and 2. In both cases, the B-section, which is not shown, differs clearly from the query melody.

For Themefinder we use the scale degrees of the incipit as query: `"1 2 3 4 5 1 6"`. Searching in all collections available results in two relevant hits at ranks 3 and 5, shown in Figure 2 as melodies R4 and R5. These are the songs *Loot ons noch ens drinken* and *Ueber die Beschwerden dieses Lebens* from the Essen collection. In melody R4, the B-section is absent, but melody R5 has a B-section that resembles the B-section of the Dutch song. The title indicates that R5 is a German translation of the original French song.

Since MELDEX also searches the Essen collection, we expect the same results. Querying with the note sequence `"g a b c' d' g' e'"` indeed results in these two hits at ranks 7 and 11. The other hits can be considered false positives.

When we search in the Musipedia database with the sequence `"g'4 a'4 b'4 c''4 d''2 g''2 e''4"`, the first hit is a song entitled *Scottish Simple de Guemene*, which is melody R6 in Figure 2. This melody is quite distant from the query. The first line ends with a half cadence and the continuation is dissimilar. Musipedia can also “search the web”. With the same query we find a nursery song called *I'm a little tea pot*, melody R7 in Figure 2. There is no B-section. The second half of the melody is like the A-section, but with a different ending.

At last, we use YahMuugle,[19] a search engine that is not designed for folk song melodies, but searches in a collection of 476,000 RISM incipits. These are incipits of classical compositions that can be found in manuscripts written before c.1800. The transportation distance algorithm described in Typke *et al.* 2003 is used. The query can be created by clicking keys on a graphical keyboard. With the sequence `"g'4 a'4 b'4 c''4 d''2 g''2 e''4"` played on the keyboard as query, we find at rank 13 melody R8 in Figure 2. According to the meta data in the result list, this is a composition by Aloys Schmitt (1788-1866) called *Variations*. Since we only get an incipit, it is not possible to compare the continuation without looking up the piece in the source manuscript.

A text-based search in the Database of Dutch Songs shows that in the Dutch language at least 30 texts exist for this melody. This finding and the results for the melody searches show that this melody must have been well known in many European countries.

Some questions concerning these search results remain to be answered. The 'skeleton' of the melody is rather generic. It consists of an embellished ascending and descending movement, which is a common contour for folk song melodies (Huron 1995). Therefore, one could argue that some of the results in Figure 2 could have been created independently instead of being derived from one original melody.

A distinctive feature of this particular melody could be the quick stepwise ascension followed by a leap to the octave and a leap back to the sixth at the very beginning. The cases where a similar B-section is present (R1 and R5) are most certainly connected. For R5 the correspondence of the German and French titles is a very strong indication. Dropping the B-section is a likely simplification of the melody. Hence, the melodies without a B-section (R4 and R7) might also have historical links with the query melody. In case of a different B-section (R2, R3 and R6), we have the least certainty of a historical link.

This extensive example shows that the currently available search engines are helpful tools for studying the history and dissemination of a tune. Although this is a quite successful example, one cannot state that the current available search functionality is sufficient for research to any folk song tune. The important feature that all these variants share is the characteristic beginning. The search engines are able to match melodies using that feature. But folk songs may have variants that are similar in other aspects, e.g., a second or third phrase may be shared while the incipits differ, or the similarity may be based upon shared melodic motifs or patterns while the global features are not particularly similar. For better results, knowledge about the process of oral transmission has to be incorporated.

4. Problems and Challenges

Based on the previous sections we can identify a number of issues that have to be addressed when taking a computational approach to study the questions of FSR.

Problem 1. There is no generally accepted theory of oral transmission of melodies. This is related to the lack of proper understanding of human musical cognitive processes such as encoding a song in human memory, performing a song from memory and creating a new song. Knowledge of these processes is an important ingredient for a theory of oral transmission of melodies.

Problem 2. Most models and concepts from FSR are not directly implementable. For implementing a model, a very precise and unambiguous description of data structures and algorithms is required. In most cases this kind of precision is not available. As Leonard Meyer states in a more general context (1996, p. 64): "[…] I have no doubt about the value of employing computers in such studies [on musical style], not merely because they can save enormous amounts of time but, equally important, because their use will force us [music scholars] to define terms and traits, classes and relationships with precision – something most of us seldom do."

Problem 3. The diversity of classification systems (see section 2.3) indicates that no universally applicable system of ordering folk song melodies exists yet. The existing systems are all designed for specific corpora. Once a researcher is familiar with a certain corpus, he might be able to determine some set of features that is expected to be discriminative for that corpus, but apparently there is no theory that for any collection can predict which features are able to be discriminative for tune families.

Problem 4. Formal testing of classification systems has not been done to a great extent. It is hard to determine whether a certain system 'works' concerning its ability to group melodies that are in the same tune family. The question is to test against what? Since the historic relationships between melodies are untraceable, it is difficult to assemble 'ground truth' data.

Not only classification systems, but also theoretic overviews such as the "Systematik" of Wiora (1941) have not been tested. He provides many examples, but no formal test that is able to convince us that his overview of melodic changes lists the changes that occur in the transmission of songs indeed.

5. Directions for Research

Based on the overviews in the previous sections, we now provide a list of open questions for future research. This section describes the research questions for possible future work, while the next section addresses the question what strategies to pursue.

The basic problem can be stated as follows: Given the availability of a symbolically encoded corpus of folk song melodies, given oral variation as characteristic feature of this corpus, given the already established approaches to categorize melodies from oral culture (section 2.3), and given the computational studies on folk songs already undertaken (section 3.1), how to develop implementable models of folk song melodies that explain the interrelationships caused by the process of oral transmission? The research questions that are relevant for this task are of interest for the MIR community, since such models could be exploited in search engines, for FSR, for Cognitive Musicology, and for Musicology in general.

5.1. Cognitive Approaches

The basic underlying problem from a cognitive point of view is: how and to what extent is a melody encoded in human memory and how is it transformed into an audible song instance during performance? This knowledge can be used to discriminate between stable and unstable elements of melodies in oral transmission. In the WITCHCRAFT project, a corpus of melodic variants that underwent this process is available to study this kind of melodic variation (Volk *et al.* 2008).

As we saw, folk song classification systems are to a certain extent based upon musical intuition of the researcher. Since musical intuition cannot be implemented, a proper understanding of it is needed, so that an implementable model can be developed. Therefore it is necessary to study this musical intuition by e.g., trying to find patterns in human descriptions of musical similarity (Volk *et al*. 2008).

The constraint-based approach proposed at the end of section 2.1 is promising for revealing the variability in memory for melodies. The challenge is to infer the constraints that led to a particular melody or a particular group of melodies from the corpus. The set of constraints thus found shows the invariable aspects of the melodies. The aspects of the melodies that are not determined by these constraints show what may vary among the instances of a certain tune. This may lead to hypotheses about the characteristics of human melodic memory, which may be tested using other kinds of experiments.

The gained insights into the cognitive aspect of melodies can be used in a more general scope of MIR tasks. Knowledge about the relation between a desired melody and the way this melody is sung from memory can be helpful in processing melodic queries for a Query By Humming system. This knowledge helps to identify the most persistent aspects of a melody and it can make recognition of a melody more robust.

An interesting property of folk song corpora from the perspective of Cognitive Musicology is that the melodies are produced by common people with normal musical skills. At least, this is the case for the collection of Dutch songs that is kept by the Meertens Institute. Regardless the quality of the melodies from an artistic point of view, all these melodies are products of some human activity of musical performance. As pointed out by Isabelle Peretz (2006, section 2) most people without a formal musical training share a ‘common core of musical knowledge’. Although vital for understanding the nature of music, the study of this common knowledge has been underrated for a long time in cognitive research. Hypotheses about common musical skills might be tested using folk song material.

5.2. The perspective of FSR

Within FSR, the research on tune families and genetic relationships of melodies seems to have been marginalized during the last two decades (Nettl 2005, p. 130). The use of new computational methods to explore and unlock collections of melodies may result in renewed interest in this topic. The enormous increase in computational power enables the development of new kinds of algorithms that incorporate more musical knowledge and that are allowed to be computationally more demanding than the ones developed during the third quarter of the twentieth century. This enables new ways to explore the contents of the many archives of folk music. This also leads to new ways to compare different corpora of melodies from different oral traditions with each other. An example of such a comparison can be found in the work of Zoltán Juhász that has been mentioned in section 3.1.2. He compares the properties of the contours of songs from various folk song traditions, which leads to hypotheses about the historical relationship of these traditions.

A common interest of FSR and MIR is the access to collections of recordings or transcriptions of folk songs. Therefore the question to be answered is: what searching or browsing functionality is needed to get access to the melodic content of a corpus of folk song melodies? More basically: what are the desired ways to get access to such a corpus from a user's point of view? How do these ways of access differ for various user groups, such as professional folk song researchers, historians, the general public, etc?

For MIR in general the question to the needs of potential users is important. To avoid mismatches between the provided tools and the demands of various groups of users, it is necessary to get a realistic overview of possible applications of MIR technology. For the particular case of Musicology, the availability of useful tools might also stimulate the acceptance of computational approaches as part of the generally available research methods.

5.3. Repeating Patterns and Stylistic Studies

The research on genetic relationships of melodies can be considered an instance of the more general musicological problem of relating instances of music to each other. In the current context the relations occur on the level of melodic contents. One approach to relate pieces of music has been proposed by Leonard Meyer (1996). His theory of musical style is founded on the replication of musical patterns. These patterns could occur on various levels of abstraction such as actual note sequences, harmonic progressions, the structure of musical compositions, etc. Instances of music that share patterns are stylistically related. Meyer proposes a hierarchy of relations, reaching from intra-opus style, e.g., the particular style of a certain symphony, via inter-opus style, e.g., the style of the oeuvre of a certain composer, to the level of the great style periods, e.g., all Baroque works. For folk songs a similar hierarchy could be conceived: at the lowest level the style of a certain song instance, then the style of a tune family, the style of all songs from a certain oral tradition, and maybe on the highest level all singable melodies. On each level the style could be characterized with the patterns that replicate in all stylistically related instances on that level. Since Meyer's approach is based on frequencies of pattern occurrences, it offers possibilities for employment in a quantitative, computational approach.

Some recent computational approaches focus on repeating patterns in melody as well, such as the theories of Ahlbäck (2004), Lartillot (2004), or the viewpoint approach of Conklin and Anagnostopoulou (2001). Also in Ethnomusicology there is interest in repeated patterns. Bruno Nettl (2005, p. 118) raises the question what are the basic units that are transmitted in an oral tradition. In his view these are musical motifs that in various recombinations form different songs. This is in accordance with the new conception of 'tune family' by Cowdery (1984), which focuses on motifs rather than solely on entire melodies. All together the detection of various kinds of repeating patterns is a promising approach to study the musical contents of a collection of melodies.

5.4. Sequence Alignment

To compare two melodies, it is insightful to notate the one below the other such that the corresponding parts are aligned. Therefore, alignments of two or more songs are heavily used by folk song researchers. The overview of results in Figure 2 is an example of such an alignment. In Computer Science alignment algorithms were developed some decades ago. Because alignments have found an application field in molecular biology, where they are used to find corresponding patterns in protein or nucleotide sequences, it is in that discipline that many algorithms, improvements and optimizations were developed, which have the potential to be employed for musicological research as well. For each alignment a score can be computed. The higher the score, the better the sequences could be aligned. In the computation of the score, domain knowledge can be incorporated. For aligning folk songs, musical scoring schemes should be developed (Van Kranenburg *et al*. 2009).

5.5. Evaluation and Testing

Finally, an important question is how to evaluate a search engine for folk song melodies. It is usually done by manually defining an ideal ranked result list for a query and comparing the results of the algorithm to that. However, this assumes that it is possible to construct such a list. In practice, this task is to a great extent subjective and based on implicit musical skills. Therefore it is not possible to create an ideal result list that can be considered *the* result list for a particular query and that can serve as ground truth. A general strategy for dealing with this problem in the context of interdisciplinary research will be presented in the next section.

6. Collaboration and Integration

The goal of developing useful software for folk song research and retrieval cannot be achieved without a profound collaboration between FSR, CM and MIR. However, the research in FSR summarized in the previous sections and the methods developed in CM and MIR to retrieve and to study folk song melodies do not indicate the actuality of such a collaboration. Since MIR and FSR have the least overlap, in this section we focus on strategies to achieve better collaboration between these two disciplines and we characterize the role of CM as mediator. The significance of such a collaboration model goes beyond the study of folk song melodies. Profound collaboration seems to be absent too in the relation between Musicology and MIR in general. Although both deal with music, there seems to be a gap in the ways of understanding it. In our opinion both disciplines suffer from this lack of mutual influence.

Characterizing the gap in an extreme way, we have on the one hand folk song researchers who lack a fundamental understanding of the possibilities and limitations of computational approaches, and on the other hand MIR researchers who do not have a professional musical knowledge framework, which causes a limited view on music and the way music functions in culture.

The existence of this gap and the focus on technical solutions prevents MIR often from being more than marginally relevant to FSR (or to Musicology in general), as for instance the problematic notion of 'ground truth' demonstrates. Sometimes it seems like MIR has a stock of so-called 'experts' from which truths can be drawn. Once provided by the expert, MIR does not go beyond this ground truth, thus making it a hermetic boundary between MIR and Musicology. The relevance to Musicology is determined by the answer to the question whether algorithms are being developed merely to *reproduce* a given 'ground truth', or to evaluate the theories that are behind that 'ground truth'. The first option seems most common, while choosing the second option will obviously lead to a better understanding of music, which in turn will lead to better approaches for music retrieval.

In a recent, related paper about Computational Ethnomusicology, Tzanetakis *et al.* (2007) observe that in MIR existing computational techniques are frequently blindly applied to musical problems, without a clear musicological goal. Therefore their first guideline for Computational Ethnomusicology is to seek active collaboration with music scholars. "Experimental results should generally be interpreted by music scholars with an understanding of the specific music(s) involved". They illustrate the guideline with some examples. Along the same lines, Cornelis *et al.* (2009) plead for a better collaboration between MIR and Ethnomusicology, which would result in more cultural independent tools for content-based MIR. In the remainder of this section we will present a more abstract model for collaboration between the involved disciplines.

Before any useful software can be developed for folk song melodies, implementable models of FSR concepts are needed. As Willard McCarty (2005, chapter 1) states in a more general discussion about the relation between Computer Science and the

Humanities, the process of modelling itself is more important than the resulting models, because it is in this process that knowledge is generated about the concepts to be modelled. Therefore, the way a model fails is more interesting than the way a model succeeds, because there lies an opportunity to improve understanding. In our case, one of the most important concepts to model is tune family.

Although the modelling is more important than the models, implementations are needed for testing and for applications such as search engines. This leads to a chain of activities that will iteratively be repeated. First, the process of creating or improving the model. Second, implementing the (adapted) model. Third, the evaluation of the implemented model, leading to improvements of the model. These activities will alternate in an iterative research process. The fourth activity is the deployment of the state-of-the-art model. This is done whenever the developed model is needed for other purposes than improving the model, such as implementing a search engine.

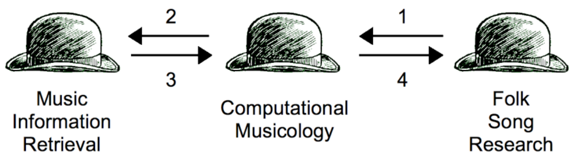

Figure 3. Three-role model for integration.

In the context of research to folk song melodies, we now present a possible way to overcome the observed ‘gap’ with the help of the three-role model that is shown in Figure 3. This model aims to identify the actors in the research process. In addition to the roles of MIR and FSR researchers, a ‘man in the middle’ role is needed. This role can be fulfilled by Computational Musicology. Since this is a role model, it does not necessarily imply the need for three separate persons in research teams. In exceptional cases one person might combine all three roles, but it would be more common for researchers either to combine both the MIR and CM or (probably more rarely) the FSR and CM roles. We now briefly characterize the three roles.

Folk Song Research. The data of Folk Song Research, such as recordings, transcriptions, notes, etc., are gathered during ethnomusicological fieldwork. Methods to process these data are necessary to support research on the relations between those artifacts and the cultural context they were obtained from. It is among these methods that classification and identification as described in section 2 can be found. For these methods a computational approach can be taken.

Since ambiguous or intuitive concepts are difficult to implement, the task for FSR is to be as precise as possible in defining concepts such that they are suited for computational modelling. Therefore, the collaboration with CM and MIR will lead to a better understanding of FSR concepts. FSR should also be involved in the evaluation of implemented models and in the process of improving these models.

In the case that the roles are not shared by the same researcher, another more practical effort is expected from folk song researchers, namely to take some time to learn how to use provided computational systems. A general understanding of the possibilities and limitations of computational methods will result in the understanding that these methods will not replace currently used methods, but open new perspectives to explore the data and to evaluate the usefulness of FSR concepts.

Music Information Retrieval. MIR designs and implements musical information systems, with the aim to improve the access to large (distributed) collections of music. Therefore, systems are generally evaluated in terms of retrieval or classification accuracy rather than by the resulting insight in the musical content. Nevertheless, as far as representations of music are used that are employable for FSR, MIR can provide numerous useful software components and user interface components that can support Folk Song Research. In our opinion MIR should employ models of music that are meaningful for FSR, and for Musicology in general. A better understanding of music will lead to more robust and more flexible music retrieval systems.

Since MIR is interested in handling large amounts of musical data, efficiency is an important constraint. Therefore, MIR could provide strategies for handling huge databases. Since CM is not primarily interested in performance issues, collaboration of MIR and CM is necessary to develop optimized algorithms that can be applied on a large scale.

Computational Musicology. In the context of this paper, the task for the CM-role is to intermediate between MIR and FSR. The subject of interest is the same as for FSR, namely songs and their relationships. The methods, however, are from Computer Science rather than from the tradition of FSR itself. These include algorithms, data structures, evaluation strategies, etc. For the activity of modelling, the task is to 'deconstruct' FSR-concepts in order to derive implementable models (arrows 1 and 2). After the first iteration these models can be improved by providing FSR the implemented models and letting FSR examine the way in which the previous models fail (arrows 3 and 4). Another, more practical, task for CM is to provide FSR with ready-to-use software frameworks and toolboxes, which allow the combining of input, processing and output methods in various ways (Garbers 2004). These toolboxes could consist of basic melodic transformations, feature extractors, segmentation algorithms, distance measures, clustering algorithms, classification methods, visualization tools, etc. that are relevant for evaluating musicological concepts. CM should hide implementation details that have no meaning in the musical domain. An example of such a toolbox is the Humdrum toolbox (Huron 1999a). However, the use of this toolbox requires a level of mastering the Unix command line environment that most musicologists do not have. The search engine Themefinder is an example of a way to provide the functionality of some Humdrum tools to users that do not have the skills to handle Humdrum directly. CM could provide interfaces like this for specific tasks.

Here FSR, MIR and CM are presented as roles rather than as disciplines to stress that it is not desirable to create a need for special CM researchers to whom FSR and MIR researchers have to obey. In practice probably all researchers within the MIR community play the role of CM to a certain extent because of their interest in music. Obviously, the quality of musical models used in music information systems affects their successfulness to a large extent. Therefore, we believe that the MIR community can gain much from pursuing the CM-role more ambitiously.

7. Concluding Remarks

The aim of FSR to identify melodies seems currently too ambitious to perform automatically (see section 2.2), since no proper implementable model of the tune family concept is available. Therefore software should support identification by finding related melodies, leaving the decision whether to assign these melodies to a specific tune family to the investigator. So, for CM, on the short term, classification tasks offer more opportunities than identification tasks. On the long term, results of that could be used to work towards a model of oral transmission. Thus CM contributes directly to the basic questions of FSR.

From the classification approaches in section 2.3, we can obtain a number of relevant features, such as cadence and accent note patterns, number of lines, and rhythmic characteristics. However, it will not be sufficient just to implement the models of e.g. Bartók or Bronson, since their feature sets were not assembled with the power of computational methods in mind, and they were fitted to specific corpora. The possibilities that Computer Science offers and the currently available computational power enable new kinds of models. Therefore, other features might be used, such as contours, repeating patterns, features from music cognition, features that reflect performance of untrained singers, and so on. Several of these features have already been used (section 3.1) or are currently being explored (Volk *et al.* 2007, Garbers *et al.* 2007). These new methods have to be developed in an interdisciplinary research context as described in section 6. Collaboration between FSR, CM and MIR provides the musical insights for computational modelling of relevant features, and for improving failing models, thus escaping the problems of ground truths that were discussed in section 6. We envision an iterative process of modelling and implementing that will result in an increasing understanding of the concepts of Folk Song Research, in particular the concept of 'tune family'. This knowledge is highly valuable for both Folk Song Research and Music Information Retrieval, and might also be of interest for other disciplines, Music Cognition in particular.

Acknowledgements

This work was supported by the Netherlands Organization for Scientific Research within the WITCHCRAFT project NWO 640-003-501, which is part of the Continuous Access to Cultural Heritage (CATCH) program.

References

Aarden, B., and D. Huron. 2001. Mapping European Folksong: Geographical Localization of Musical Features. *Computing in Musicology* 12: 169-183.

Ahlbäck, S. 2004. *Melody beyond Notes: A Study of Melody Cognition*. PhD diss. Göteborgs Universitet.

Bayard, S.P. 1950. Prolegomena to a Study of the Principal Melodic Families of British-American Folk Song. *Journal of American Folklore* 63(247): 1-44.

Bod, R. 2002. Memory-Based Models of Melodic Analysis: Challenging the Gestalt Principles. *Journal of New Music Research* 31(1): 27-37.

Bohlman, P.V. 1988. *The Study of Folk Music in the Modern World*. Bloomington: Indiana University Press.

Brandts Buys, M.A. 1975. *Gezelschapsliederen*. Leiden: A.W. Sijthoff.

Bronson, B.H. 1949. Mechanical Help in the Study of Folk Song. *The Journal of American Folklore* 62(244): 81-86.

Bronson, B.H. 1950. Some Observations about Melodic Variation in British-American Folk Tunes. *Journal of the American Musicological Society* 3: 120-134.

Bronson, B.H. 1951. Melodic Stability in Oral Transmission. *Journal of the International Folk Music Council* 3: 50-55.

Bronson, B.H. 1959. Toward the Comparative Analysis of British-American Folk Tunes. *The Journal of American Folklore* 87(284): 165-191.

Clarke, E., and N. Cook (eds.). 2004. *Empirical Musicology: Aims, Methods, Prospects*. Oxford: Oxford University Press.

Conklin, D. 2006. Melodic Analysis with Segment Classes. *Machine Learning* 65(2-3): 349-360.

Conklin, D., and C. Anagnostopoulou. 2001. Representation and Discovery of Multiple Viewpoint Patterns. Paper presented at the International Computer Music Conference, September 17-22, in Havana, Cuba.

Cornelis, O., M. Lesaffre, D. Moelants, and M. Leman. 2009. Access to ethnic music: Advances and perspectives in content-based music information retrieval. Signal Processing. doi:10.1016/j.sigpro.2009.06.020. [Article in Press].

Cowdery, J.R. 1984. A Fresh Look at the Concept of Tune Family. *Ethnomusicology* 28(3): 495-504.

Eck, D., and N. Casagrande. 2005. Finding Meter in Music Using an Autocorrelation Phase Matrix and Shannon Entropy. Paper presented at the 6th International Conference on Music Information Retrieval (ISMIR 2005), September 11-15, in London, UK.

Eerola, T., and P. Toiviainen. 2004. MIR in Matlab: The MIDI Toolbox. Paper presented at the 5th International Conference on Music Information Retrieval (ISMIR 2004), October 10-24, in Barcelona, Spain.

Elbourne, R.P. 1975. The Question of Definition. *Yearbook of the International Folk Music Council* 7: 9-29.

Elscheková, A. 1966. Methods of Classification of Folk-Tunes. *Journal of the International Folk Music Council* 18: 56-76.

Frieler, K. 2007. Visualizing Music on the Metrical Circle. Paper presented at the 8th International Conference on Music Information Retrieval (ISMIR 2007), September 23-27, in Vienna, Austria.

Garbers, J. 2004. *Integration von Bedien- und Programmiersprachen am Beispiel von OpenMusic, Humdrum und Rubato*. PhD diss. Fakultät IV – Elektrotechnik und Informatik der Technischen Universität Berlin, http://people.cs.uu.nl/garbers/publications/DissertationJoergGarbers.pdf (accessed July 2, 2008).

Garbers, J., P. van Kranenburg, A. Volk, F. Wiering, R.C. Veltkamp, and L.P. Grijp. 2007. On Pitch and Chord Stability in Folk Song Variation Retrieval. Paper presented at the First International Conference of the Society for Mathematics and Computation in Music (MCM 2007), in Berlin, Germany.

Hoos, H.H., K. Renz, and M. Görg. 2001. GUIDO/MIR: An Experimental Musical Information Retrieval System Based on GUIDO Music Notation. Paper presented at the 2nd International Symposium on Music Information Retrieval (MUSIC IR 2001), October 15-17, in Bloomington (IN), USA.

Huron, D. 1995. The Melodic Arch in Western Folksongs. *Computing in Musicology* 10: 3-23.

Huron, D. 1999a. *Music Research Using Humdrum: A User's Guide*. Stanford (CA): Center for Computer Assisted Research in the Humanities, http://csml.som.ohio-state.edu/Humdrum/guide.toc.html (accessed August 27, 2007).

Huron, D. 1999b. *The New Empiricism: Systematic Musicology in a Postmodern Age*. Lecture 3 from the 1999 Ernest Bloch Lectures, http://musiccog.ohio-state.edu/Music220/Bloch.lectures/3.Methodology.html (accessed July 2, 2008).

Jesser, B. 1991. *Interaktive Melodieanalyse*. Bern: Peter Lang.

Juhász, Z. 2000. Contour Analysis of Hungarian Folk Music in a Multidimensional Metric-Space. *Journal of New Music Research* 29(1): 71-83.

Juhász, Z. 2002. The Structure of an Oral Tradition: Mapping of Hungarian Folk Music to a Metric Space. *Journal of New Music Research* 31(4): 295-310.

Juhász, Z., 2004. Segmentation of Hungarian Folk Songs Using an Entropy-Based Learning System. *Journal of New Music Research* 33(1): 5-15.

Juhász, Z. 2006. A Systematic Comparison of Different European Folk Music Traditions Using Self-Organizing Maps. *Journal of New Music Research* 35(2): 95-112.

Keller, M.S. 1984. The Problem of Classification in Folksong Research: A Short History. *Folklore* 95(1): 100-104.

Klusen, E., H. Moog and W. Piel. 1978. Experimente zur mündlichen Tradition von Melodien. *Jahrbuch für Volksliedforschung* 23: 11-32.

Krohn, I. 1903. Welche ist die beste Methode, um Volks- und volksmäßige Lieder nach ihrer melodischen (nicht textlichen) Beschaffenheit lexikalisch zu ordnen? *Sammelbände der Internationalen Musikgesellschaft* 4, H. 4: 643-660.

Lartillot, O. 2004. A Musical Pattern Discovery System Founded on a Modeling of Listening Strategies. *Computer Music Journal* 28(3): 53-67.

McCarty, W. 2005. *Humanities Computing*. Basingstoke: Palgrave Macmillan.

McNab, R.J., L.A. Smith, D. Bainbridge, and I.H. Witten. 1997. The New Zealand Digital Library MELody inDEX. *D-Lib Magazine* May, http://www.dlib.org/dlib/may97/meldex/05witten.html (accessed July 2, 2008).

Meyer, L.B. 1996. *Style and Music – Theory, History, and Ideology*. Chicago: University of Chicago Press.

Nettl, B. 2005. *The Study of Ethnomusicology*. 2nd edition. Urbana: University of Illinois Press.

Pearce, M.T., D. Müllensiefen, and G.A. Wiggins. 2008. A Comparison of Statistical and Rule-Based Methods of Melodic Segmentation. Paper presented at the 9th International Conference on Music Information Retrieval, September 14-18, in Philadelphia (PA), USA.

Peretz, I. 2006. The Nature of Music from a Biological Perspective. *Cognition* 100: 1-32.

Peretz, I., and R.J. Zatore. 2005. Brain Organization for Music Processing. *Annual Review of Psychology* 56: 89-114.

Pickens, J. 2000. A Comparison of Language Modeling and Probabilistic Text Information Retrieval Approaches to Monophonic Music Retrieval. Paper presented at the 1st International Symposium on Music Information Retrieval (MUSIC IR 2000), October 23-25, in Plymouth, Massachusetts.

Rubin, D.C. 1995. *Memory in Oral Traditions*. New York: Oxford University Press.
Sapp, C. 2005. Online Database of Scores in the Humdrum File Format. Paper presented at the 6th International Conference on Music Information Retrieval (ISMIR 2005), September 11-15, in London, UK.
Schaffrath, H. 1992. The Retrieval of Monophonic Melodies and Their Variants – Concepts and Strategies for Computer-Aided Analysis. In: A. Marsden and A. Pople (eds.). *Computer Representations and Models in Music*. London: Academic Press. 95-109.
Schaffrath, H. 1995. *The Essen Folksong Collection*. D. Huron, ed. Stanford (CA): Center for Computer Assisted Research in the Humanities.
Schaffrath, H. 1997. The Essen Associative Code: A Code for Folksong Analysis. in: E. Selfridge-Field (ed.). *Beyond Midi – The Handbook of Musical Codes*. Cambridge: The MIT Press. 343-361.
Singer, J. 2004. Creating a Nested Melodic Representation: Competition and Cooperation among Bottom-Up and Top-Down Gestalt Principles. Paper presented at the 5th International Conference on Music Information Retrieval (ISMIR 2004), October 10-15, in Barcelona, Spain.
Skalak, M., J. Han, and B. Pardo. 2008. Speeding Melody Search with Vantage Point Trees. Paper presented at the 9th International Conference on Music Information Retrieval (ISMIR 2008), September 14-18, in Philadelphia (PA), USA.
Stein, E. 1979. The Use of Computers in Folklore and Folk Music: A Preliminary Bibliography. The American Folklife Center, http://www.loc.gov/folklife/guides/BibComputers.html (accessed July 2, 2008).
Steinbeck, W. 1982. *Struktur und Ähnlichkeit: Methoden automatisierter Melodieanalyse*. Kassel: Bärenreiter.
Suchoff, B. 1981. Preface to *The Hungarian Folk Song*, by B. Bartók, ix-lv. Albany: State University of New York Press.
Suppan, W., and W. Stief (eds.). 1976. *Melodietypen des Deutschen Volksgesanges*. Tutzing: Hans Schneider.
Temperley, D. 2006. A Probabilistic Model of Melody Perception. Paper presented at the 7th International Conference on Music Information Retrieval (ISMIR 2006), October 8-12, in Victoria (BC), Canada.
Toiviainen, P., and T. Eerola. 2005. Classification of Musical Metre with Autocorrelation and Discriminant Functions. Paper presented at the 6th International Conference on Music Information Retrieval (ISMIR 2005), September 11-15, in London, UK.
Typke, R., P. Giannopoulos, R.C. Veltkamp, F. Wiering, and R. van Oostrum. 2003. Using Transportation Distances for Measuring Melodic Similarity. Paper presented at the 4th International Conference on Music Information Retrieval (ISMIR 2003), October 26-30, in Baltimore (MD), USA.
Tzanetakis, G., A. Kapur, W.A. Schloss, and M. Wright. 2007. Computational Ethnomusicology. *Journal of interdisciplinary music studies* 1(2): 1-24.
Van Kranenburg, P., A. Volk, F. Wiering, and R.C. Veltkamp. 2009. Musical Models for Folk-Song Melody Alignment. Paper presented at the 10th International Conference on Music Information Retrieval (ISMIR 2009), October 26-30, in Kobe, Japan.
Volk, A., J. Garbers, P. van Kranenburg, F. Wiering, L.P. Grijp, and R.C. Veltkamp. 2007. Comparing Computational Approaches to Rhythmic and Melodic Similarity in Folk Song Research. Paper presented at the First International Conference of the Society for Mathematics and Computation in Music (MCM 2007), in Berlin, Germany.
Wiering, F., R.C. Veltkamp, J. Garbers, A. Volk, P. van Kranenburg, and L.P. Grijp. 2009. Modelling Folksong Melodies. *Interdisciplinary Science Reviews* 34(2–3): 154–171.
Wiora, W. 1941. Systematik der musikalischen Erscheinungen des Umsingens. *Jahrbuch für Volksliedforschung* 7: 128-195.

[1] The definition of the term 'folk song' is not without problems. One of the most stable ingredients in the many attempts to define the concept is the process of oral transmission. "For an item to qualify as folklore it must have been in oral circulation, passing from individual to individual without aid of any written text." (Elbourne 1975).
[2]Some projects that illustrate this priority are the EthnoArc project (http:// http://www.ethnoarc.org, accessed 20 October 2008), the DISMARC project (http://www.dismarc.org, accessed 20 October 2008) and the WITCHCRAFT project (http://www.cs.uu.nl/research/projects/witchcraft, accessed 20 October 2008).
[3]Acronym for: 'What Is Topical in Cultural Heritage: Content-based Retrieval Among Folksong Tunes'.
[4] Nederlandse Liederenbank, http://www. liederenbank.nl (accessed 12 September 2007).
[5] Everything in a melody can change.
[6] This causes an ambiguity in the term 'song', with which an individual performance can be meant, but also the tune family as a whole.
[7] The meaning of the term 'classification' in the context of FSR is not equivalent to the use of the term in Computer Science areas such as Pattern Recognition and Machine Learning. In Computer Science the relation between classes is not necessarily defined, while in FSR the classes are put into some kind of rational order as well, e.g., first all songs consisting of one phrase, then all songs consisting of two phrases, and so on.
[8] Juhász uses the term 'section' rather than 'phrase' or 'line'. This probably finds its origin in the terminology used by Béla Bartók.
[9] Eck and Casagrande 2005, Eerola and Toiviainen 2004, Frieler 2007, Hoos *et al.* 2001, Pearce *et al.* 2008, Pickens 2000, Sapp 2005, Singer 2004, Skalak *et al.* 2008, Temperley 2006, Toiviainen and Eerola 2005.
[10] http://www.dafos.dk/melodies-online.aspx (accessed 2 July 2008).
[11]http://www.dafos.dk/melodies-online/melody-codes-and-code-searching.aspx (accessed 2 July 2008).
[12] http://www.colonialdancing.org/Easmes (accessed 1 October 2008).
[13] http://esavelmat.jyu.fi (accessed 1 October 2008).
[14] http://www.themefinder.org (accessed 2 July 2008).
[15] http://www.nzdl.org/musiclib (accessed 2 July 2008).
[16] http://www.mudcat.org (accessed 2 July 2008).
[17] http://www.musipedia.org (accessed 2 July 2008).
[18] http://www.liederenbank.nl (accessed 2 July 2008).
[19] http://www.yahmuugle.cs.uu.nl (accessed 3 July 2008).

Biographies

Peter van Kranenburg obtained masters degrees in electrical engineering (2003, Delft University of Technology) and Musicology (2004, Utrecht University). In his master studies he developed machine learning methods for studying authorship of musical compositions, with a particular application to authorship problems of several organ fugues ascribed to J.S. Bach. Since 2006, he is Ph.D. researcher at Utrecht University in the WITCHCRAFT project. His research focuses on the computational modeling of similarity of folk-song melodies for retrieval purposes.

Jörg Garbers is a computer scientist with many years of project experience in computational linguistics and computational musicology. As a leading software developer, he took care of the development and interoperability of systems for music analysis, composition and retrieval. In the vision underlying his Ph.D. thesis, technically informed musicologists and creative users explore, modify and combine each others software components and thereby collaboratively develop improved musical models and computational methods.

Anja Volk holds masters degrees in musicology and mathematics. In 2002 she received a Ph.D. in the area of computational musicology from Humboldt University Berlin. She has worked in the area of Mathematical Music Theory at the Technical University of Berlin (1998-2003), in the area of Computation and Cognition in Music at the University of Southern California (2003-2005), and is currently working in the area of Music Information Retrieval applied to Folksong research in the WITCHCRAFT-project at Utrecht University and the Meertens Institute in Amsterdam.

Frans Wiering received a Ph.D. in Musicology (1995, University of Amsterdam) for his thesis *The Language of the Modes. Studies in the History of Polyphonic Modality*. Since the late 1980s, he has been researching computer applications in musicology. He is the founder of the Thesaurus musicarum italicarum (www.euromusicology.org). His present research is in digital scholarly publishing of music and music information retrieval. He was a Visiting Scholar at Stanford University and a Visiting Fellow at Goldsmiths College, University of London. He is currently an Assistant Professor at the Department of Information and Computing Sciences of Utrecht University (Netherlands).

Louis Peter Grijp (The Hague 1954) studied musicology at Utrecht University and lute at the Royal Conservatory of The Hague. His dissertation (Utrecht 1991) treated the mechanism of contrafact writing in Dutch songs of the seventeenth century. Dutch songs became Grijp's specialism as research fellow of the Meertens Institute in Amsterdam and as professor of Dutch song culture at Utrecht University. Under his supervision the Nederlandse liederenbank (Dutch Song Database) has been developed in the Meertens Institute. Grijp is member of the Royal Dutch Academy of Arts and Sciences and received several prizes for both his scholarly and musical work.

Remco Veltkamp studied computer science at Leiden University and completed his M.Sc. thesis at IBM Scientific Centre, Paris, France. He received the Ph.D. degree from Erasmus University, Rotterdam, in 1992. Since 1995 he has worked at Utrecht University, where he contributes to the master program Game and Media Technology. He is full professor in Multimedia. His current research focuses on the algorithmic aspects of media technology, like the algorithm design and analysis, and experimental verification. He is Editor of the international journal Pattern Recognition and the International Journal on Shape Modeling.