Towards Integration of Music Information Retrieval and Folk Song Research

Peter van Kranenburg, Jörg Garbers, Anja Volk, Frans Wiering, Louis Grijp, Remco C. Veltkamp

Department of Information and Computing Sciences, Utrecht University

Technical Report UU-CS-2007-016

www.cs.uu.nl

ISSN: 0924-3275

Abstract

Folk song research often deals with large collections of tunes that have various types of relations to each other. Computational methods can support the study of the contents of these collections. Music Information Retrieval research provides such methods. Yet a fruitful cooperation of both disciplines is difficult to achieve. We give overviews of relevant work from both disciplines, and we provide research directions and a collaboration model that may lead to better cooperation.

1 Introduction

The goal of the WITCHCRAFT project ('What Is Topical in Cultural Heritage: Content-based Retrieval Among Folksong Tunes') 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 [40]. 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.

'Folk songs' were sung by common people during work or social activities. One of their most important characteristics is that they are part of oral culture.[1] The melodies and the texts are learned by imitation and participation rather than from books. In the course of this oral transmission, changes occur to the melodies. The resulting set of variants of a song forms a so called 'tune family'. Many folk songs were recorded on tape or transcribed on paper by field workers and are thus available for research.

Although attention has been paid to folk songs in the Music Information Retrieval (MIR) community,[2] 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. But 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-formalised nature of many FSR concepts and theories. Therefore we need to find an approach to bridge this gap. In this paper we first give overviews of relevant work that has been done in both disciplines, and after that we provide directions for research and we describe an approach that may lead to better cooperation.

2 Folk Songs in MIR

Only a limited number of MIR applications and studies are specifically aimed at searching folk song collections. In this section a number of search engines and research papers will be surveyed.

[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.' [17].

[2]In this paper 'MIR' is taken in its broadest sense: not only specific retrieval research, but also other research that has computational processing of music as its object.

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 [13] contains about 10,000 instrumental melodies found in books and manuscripts with Danish folk music. 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 bars), corresponding to the 'beats' 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 [14]. 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.

Another online searchable database with folk song tunes is the Digital Archive of Finnish Folk Tunes [20]. The collection contains about 9,000 tunes, most of them collected in the early 20th century by Ilmari Krohn. Two types of 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. Wild-cards 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 tunes.

Themefinder [54] can be used to search a large collection of themes and incipits from classical works, folk songs (the Essen collection [45]) and 16th-century Latin motets. 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 Meter. The user can choose whether the query should occur at the beginning, or anywhere in the theme. Internally the Humdrum Toolkit [24] 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 [35, 36]. 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 [37]. 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 [35]. 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 [38] contains a large number of folk song melodies. Various input methods can be used to create a query: Lilypond code, contour (Parsons code), humming/whistling, or tapping. For matching, transportations distances [56] 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.

Only the Danish search engine possesses a query method that is motivated by FSR. One can search for a sequence of accented notes, which are assumed to be rather stable across variants of a melody.

Figure 1: The Dutch song *Dat gaat naar Den Bosch toe* as notated in [6, p.81]. For the first phrase accented notes on various levels are marked.

2.2 An Example

The melody of the Dutch folk song *Dat gaat naar Den Bosch toe* 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) [40]. The first two phrases of Gaveaux' song are shown as melodies **G1** and **G2** in Figure 2. As indicated, the Dutch melody consists of two sections, A and B. Thus, the structure of the song is ABABA. As an example to show what can be accomplished by using the search engines described in the previous section, we will take this melody as a query and report about the results.

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 tune 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 meter. 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 meter 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 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.

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.

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 song.

Since Meldex also searches in 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 [59], a search engine that is not designed for folk song melodies, but searches in a collection of RISM incipits. These are incipits of classical compositions that can be found in manuscripts written before 1800. The same transportation distance algorithm as in Musipedia 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 [25]. Therefore, one could argue that 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 might 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 an 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 melody. 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 motives or patterns while the global features are not particularly similar. For better results, knowledge about the process of oral transmission has to be incorporated.

2.3 Research Papers

Folk song melodies have been used as data 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 eight papers in the complete ISMIR proceedings from 2001–2006 that employ the Essen folk song collection [15, 16, 23, 43, 46, 48, 53, 55]. In none of these papers the implications of the choice of this data set is 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 meter classification, benchmark establishing or segmentation, have not been interpreted concerning their potential to contribute to folk song research.

The studies that are most interesting for our current purpose are those that are driven by an interest in folk music as such. In the remainder of this section an overview is given of relevant studies.

2.3.1 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 [7]. Thus songs with certain desired characteristics could be retrieved using a sorting machine [9]. In the following decades many studies on the use of computers in folklore and folk music has been published. A bibliography from 1979 on this topic lists more than 100 references [49].

In the 1980s Wolfram Steinbeck [50] and Barbara Jesser [26] 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, ambitus, 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 not in all families the same features were important.

Wolfram Steinbeck focused primary on clustering [50]. He used hierarchical clustering algorithms to group melodies from the Essen collection together using 13 features such as mean and standard deviation of pitches, ambitus, size of intervals, number of direction changes, and others [50, p. 275]. With a set of 35 melodies he was able to cluster melodies into meaningful groups, such as hymns, children songs and hunting songs. An experiment with 500 melodies also led to clusters, but in this case the clusters were more difficult to characterize.

2.3.2 Contour Based Approaches

There are some contour-based approaches. In a series of experiments, David Huron confirmed the hypothesis that folk song melodies tend to show arch-like contours in single lines as well as in successions of lines using tools from his Humdrum Toolkit [25, 24]. His testing material consisted of the Essen collection.

Zoltán Juhász published several studies in which Hungarian folk song melodies are represented by contour vectors [27, 28, 29, 30]. A contour vector 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 the melodic phrases.[3] The principal components can be interpreted as contours. 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 [28]. Juhász compared the Hungarian melodies with songs from other countries in Europe by training Self Organizing Maps with the contour vectors [30]. 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 causes Juhász to speculate about a 'common language' that reflects an archaic common origin of these European traditions.

2.3.3 Segmentation

Some segmentation algorithms that have been developed within MIR were aimed for or tested on folk song melodies. Zoltán Juhász used the entropy of the continuation of a sequence of intervals [29]. A high value of this conditional entropy implies that the next interval is hard to predict, which may indicate a segment boundary. This can be used for segmentation.

Another data-driven approach to segmentation is presented by Rens Bod [4]. 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.

2.3.4 Other Approaches

Darrel Conklin and Christina Anagnostopoulou [10] developed the concept of viewpoint. This is a formalized way to represent a feature of a melody, resulting in a sequence of symbols. Using a contour viewpoint and association rule mining Conklin [11] got the same arch shape Huron found [25]. In combination with suffix trees, viewpoints can also be employed to find repeating patterns.

Bret Aarden and David Huron proposed the use of geographical information [1]. 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. However, they are very few in number compared to the entire field of MIR. Only those of Juhász explicitly state an interest in folk songs as part of oral traditions.

3 Some Past and Current Approaches to Melody in Folk Song Research

During the last century, the availability of collected folk song tunes has generated a considerable amount of musicological research. One of the primary concerns is how to deal with the specific type of variation caused by the process of oral transmission. Therefore we will first discuss oral transmission. Then classification and identification of melodies in the context of FSR will be discussed.

[3] Juhász uses the term 'section' rather than 'phrase' or 'line'. This probably finds its origin in the terminology used by Béla Bartók.

3.1 Oral Transmission

The transmission of songs in an oral tradition is determined by the capabilities of human perception, memory, performance and creativity. Participants in a tradition have 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 [41] shows that the representation of a song in human memory is not 'literal'. It is — using the words of Bertrand Bronson — 'a fluid idea of a song' [8]. During performance the actual appearance of the song is recreated. In the process of transformation from the memory representation to audible words and melody, considerable variation may occur. As long as the process of performing a song from memory is not sufficiently understood, we have to focus mainly on the recorded song instances in order to infer knowledge about this kind of variation.

A comprehensive inventory of types of variation in German folk songs is made by Walter Wiora [58], who summarizes the issue as follows: 'Alles an der Beschaffenheit einer Melodie ist veränderlich'.[4] 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, 7. demolition of the melody. He provides many examples.

To show that a 'genetical' relation does not necessarily imply a relation of musical similarity, Bruno Nettl describes a way in which a song may change entirely [39, p. 116]. 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.

After studying traditional music of Ireland, James Cowdery proposed another view on the concept of tune family [12]. In the original approach Bayard defined the concept of tune family by reference to both empirical data (melodies gathered during field work) and a hypothesized historical sequence, assuming that each melody in a tune family is a derivative of a single air [3]. Therefore, Bayard used similarities in overall contour to relate tunes that are in a tune family. Cowdery criticizes this view by posing that folk musicians do not relate new tunes to abstract archetypical melodies, but to other concrete tunes they know. Therefore, Cowdery does not only focus on global similarity but also on motifs that are shared among members of a tune family. Tunes may have some sections in common while other sections differ. In his view all tunes in a tune family are composed from melodic material from a common 'melodic pool'. This relieves us from the hypothetical historical sequence of tunes that in virtually all cases has been lost, if existed at all. Instead, all 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. This view also accounts for related melodies that do not share global characteristics, but share musical motifs.

There is no generally accepted theory that explains melodic variation in oral cultures yet. David Rubin has developed a theory for oral transmission of texts in which variation is modeled by constraint based reconstruction of the texts from memory [44]. 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 would be to model these constraints.

[4] Everything in a melody can change.

3.2 Classification

A 'classification system' is used to group together melodies that share certain characteristics. Examples are a common number of lines, a common number of syllables or a common cadence note sequence. Overviews of classification systems are provided in [18, 31, 5].

Classification is necessary for e.g. storing melodies in a card file database or for publishing a book with melodies. In those cases a linear 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. In the following we give examples for features used in different classification systems.

Within FSR, there is not one universally applicable classification system in existence. Most systems were developed for specific corpora. One of the first was developed for Finnish songs by Ilmari Krohn. In his system the cadence notes (ending notes of the lines, as depicted in Figure 1) are most important [19]. Béla Bartók and Zoltán Kodály adapted his system for Hungarian folk songs. In their publications songs were ordered by: 1. the number of lines, 2. the sequence of cadence notes, 3. the number of syllables in each line, and 4. the ambitus [51, 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 [51, 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 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 [39, p. 124]. This applies to folk song classification systems in general [5, p. 33].

For the British-American folk song tradition, Bertrand Bronson ordered a list of features according to importance using his 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. [8]. Thus, a classification system based upon these features can be expected to group songs in the same tune family together.

In an important publication of The German Archive of Folk Song (Deutsches Volksliedarchiv) containing German ballads, an ordering is used that reflects the system of Krohn [52]. The first criterion is the number of lines and the second criterion is the cadence note sequences. If a resulting group is judged incoherent by the editors, subgroups are made. The exact way in which these subgroups were established, is not accounted for.

3.3 Identification

'Identification' of a song is more directly related to the process of oral transmission. If two song instances are derived from the same common 'ancestor', they belong to the same 'tune family' [3] and are considered to be the same song [39, p. 114].[5] The identity of a song is a complex and abstract concept. It is not obvious what constitutes the 'substance' of a song that is shared among all historically derived variants.[6] As a consequence, historically linked variants may in a classification system end up in entirely different classes. The possibility of interference between tune families complicates the issue even further. Because the concept

[5] This causes an ambiguity in the term 'song', with which an individual performance can be meant, but also the tune family as a whole.

[6] This 'substance' is called 'content' by Bruno Nettl [39, p. 54f] as opposed to 'style'. The style may change, but as lang as the content is retained, it is the same song. Nettl leaves the way to separate style from contents implicit.

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?

At the Meertens Institute, the concept of 'melodienorm' (melody norm)[7] is used to group 'genetically' related melodies. Because the contents of folk song collections are highly fragmentary, it is impossible to reconstruct the complete history of melodies and to find all variants that are derived from a common 'ancestor' melody. What is feasible is 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 melody norms in a second stage. For this a retrieval system is an important tool.

Within the specific context of the Meertens Institute, identification of melodies is also important for more practical reasons. A number of the songs in the collection of the institute will appear in print, and in the near future a cd-box will be published with recordings of Dutch folk songs. For publications like these, it is important to know which melodies are in the collection and what their relations and identities are. It is unfeasible to find all relations by hand, therefore a search engine is needed.

4 Directions for Research

The basic question for the WITCHCRAFT project can be described as follows. Given the availability of a symbolically encoded corpus of Dutch folk song melodies, given oral variation as characteristic feature of this corpus, given the already established approaches to categorize melodies from oral culture (section 3), and given the computational studies on folk songs already undertaken (section 2), how to develop a search engine that is able to retrieve variants of a certain melody? Are the concepts and methods that are reviewed in sections 2 and 3 sufficient, or do we have to take new approaches? The research questions that are relevant for this task are of interest not only for the WITCHCRAFT project, but also for the MIR community as a whole, for FSR, for Cognitive Musicology, and for Musicology in general. In the remainder of this section some relevant directions for research will shortly be mentioned.

One of the basic underlying problems from a cognitive point of view is: how is a melody encoded in human memory and how is it being transformed into an audible song instance during performance? Knowing this might enable us to discriminate between stable and unstable elements of melodies in oral transmission. The availability of a corpus with melodic variants that underwent this process provides data to study this kind of melodic variation.

The constraint based approach proposed in section 3.1 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 thus found set of constraints 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 might lead to hypotheses about the characteristics of human melodic memory, which might be tested using other kinds of experiments.

[7] Comparable to 'tune family' and 'Melodietyp'.

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. Thus the research question how to construct meaningful queries can be addressed.

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 [42, 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 to 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 collected by field workers.

Within FSR, the research on tune families and genetic relationships of melodies seems to have been marginalized during the last two decades [39, 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 2.3. 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.

One of the fields of interest of FSR is the preservation of collections of songs as recordings or transcriptions. Many archives preserve collections of recordings of all kinds of music, of which the Meertens Institute is just one. The activities of the WITCHCRAFT project aim at unlocking such a collection on the basis of melodic content. 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 users point of view? And how do these ways of access differ for various user groups, such as professional folk song researchers, historians, the general public, etc? These questions will be addressed in the design of the search engine and in the integration of the search engine with the Database of Dutch Songs.

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. The context of the WITCHCRAFT project provides an opportunity to adapt existing MIR tools and to develop new tools in collaboration with folk song researchers, thus addressing the specific needs of FSR.

The research on genetic relationships of tunes 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 [33]. 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 also focus on repeating patterns in melody, such as the theories of Ahlbäck [2], Lartillot [32], or the viewpoint approach of Conklin and Anagnostopoulou [10]. Also in Ethnomusicology there is interest in repeated patterns. Bruno Nettl raises the question what are the basic units that are transmitted in an oral tradition [39, p. 118]. In his view these are musical motives that in various recombinations form different songs. This is in accordance with the new conception of 'tune family' by Cowdery [12], which focuses on motifs rather than 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.

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 for a great part 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 way of dealing with this problem will be presented in the next section.

5 Cooperation and integration

The goal of developing useful software for folk song research cannot be achieved without a profound cooperation between FSR and MIR. However, the research questions in FSR summarized in the previous sections and the methods developed in MIR to retrieve and to study melodies do not indicate the actuality of such a cooperation. This seems to be true too for 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 1. folk song researchers who lack a fundamental understanding of the possibilities and limitations of computational approaches, and 2. 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.

This situation 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.

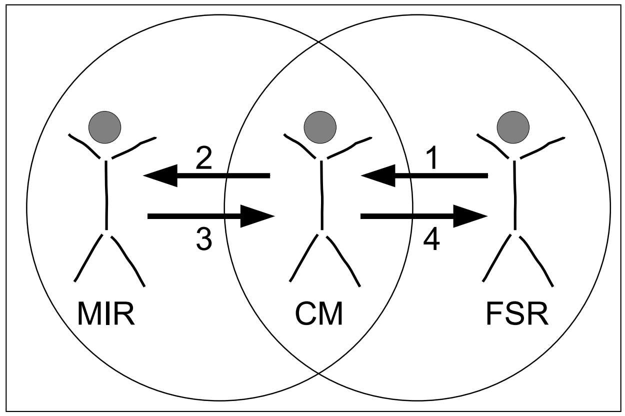

Figure 3: Three-role model for integration: Music Information Retrieval (MIR), Computational Musicology (CM), and Folk Song Research (FSR).

Once provided by the expert, MIR does not go beyond this ground truth, thus making it a hermetic boundary of MIR and Musicology. The relevance to Musicology is determined by 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 second option will obviously lead to a better understanding of music, which in turn will lead to better approaches for music retrieval.

Before any useful software can be developed for folk song research—which is a core activity within the WITCHCRAFT project—implementable models of FSR concepts are needed. As Willard McCarty states in a more general discussion about the relation between Computer Science and the Humanities [34, Ch.1], the process of modeling itself is more important than the resulting models, because it is in this process that knowledge is generated about the concepts to be modeled. 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 the melody norm, or tune family.

Although the modeling 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. And third, the evaluation of the implemented model, leading to improvements of the model. These activities will alternate in an iterative research process. A 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.

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. In addition to the roles of MIR and FSR researchers, a 'man in the middle' role is needed. We call this role 'Computational Musicology' (CM). It does not necessarily imply the need for an extra person 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.

FSR. The data of Folk Song Research, such as recordings, transcriptions, notes, etc., are gathered during field work. Therefore, methods to sort out 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 3 can be found. For these methods a computational approach might be taken.

Since ambiguous or intuitive concepts are difficult to implement, the cooperation with Computer Science offers the opportunity to gain more clear understanding of the underlying concepts. Hence, the task for FSR is to be as precise as possible in defining concepts such that they are suited for computational modeling. FSR should also be involved in the evaluation of implemented models and in the process of improving these models. Another more practical effort that may be expected from folk song researchers, is to take some time to learn how to use provided computational systems.

If folk song researchers put effort in getting a general understanding of the possibilities and limitations of computational methods, they will realize that these methods will not replace currently used methods, but provide additional ways to explore the data and to evaluate the usefulness of their concepts.

MIR. 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.

CM. In the context of this paper, the task for the CM-role is to connect 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. For the activity of modeling 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, allowing to combine input, processing and output methods in various ways [21]. 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 [24]. 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. If they were not interested in music, they would probably not have been involved in MIR at all. To model music for use in music information systems, at least some knowledge about the nature of music is needed. Therefore, the MIR community can gain much from pursuing the CM-role more ambitiously.

6 Concluding Remarks

Facing the tasks described in the previous section, what could CM learn from the current situations of MIR and FSR as described in sections 2 and 3? The aim of FSR to identify melodies seems currently too ambitious to perform automatically (see section 3.3), since no proper implementable model of melody norm 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.

From the classification approaches in section 3.2 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 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 2) or are currently being explored [57, 22]. These new methods have to be developed in an interdisciplinary research context as described in section 5. Cooperation with (computational) musicologists provides the musical insights for modeling the features, and for improving failing models, thus escaping the problems of ground truths that were discussed in section 5. We envision an iterative process of modeling and implementing that will result in an increasing understanding of the concepts of Folk Song Research, in particular the identity of a 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 CATCH-program.

References

[1] Aarden, B., and D. Huron, "Mapping European Folksong: Geographical Localization of Musical Features", *Computing in Musicology* 12 (2001), 169–183.

[2] Ahlbäck, S., *Melody Beyond Notes: A Study of Melody Cognition*, [Diss.], Göteborgs Universitet, 2004.

[3] Bayard, S.P., "Prolegomena to a Study of the Principal Melodic Families of British-American Folk Song", *Journal of Americal Folklore* 63 (1950), no. 247, 1–44.

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

[5] Bohlman, P. V., *The Study of Folk Music in the Modern World*, Bloomington, 1988.

[6] Brandts Buys, M.A., *Gezelschapsliederen*, Leiden, 1875.

[7] Bronson, B. H., "Mechanical help in the study of folk song", *The Journal of American Folklore* 62 (1949), No. 244, 81–86.

[8] Bronson, B. H., "Some Observations About Melodic Variation in British-American Folk Tunes", *Journal of the American Musicological Society* 3 (1950), 120–134.

[9] Bronson, B. H., "Toward the Comparative Analysis of British-American Folk Tunes", *The Journal of American Folklore* 87 (1959), No. 284, 165–191.

[10] Conklin, D., and C. Anagnostopoulou, "Representation and Discovery of Multiple Viewpoint Patterns", *Proceedings of the 2001 International Computer Music Conference*, Havana, Cuba, 2001, 479–485.

[11] Conklin, D., "Melodic analysis with segment classes", *Machine Learning* 65 (2006), nr. 2-3, 349–360.

[12] Cowdery, J.R., "A Fresh Look at the Concept of Tune Family", *Ethnomusicology* 28 (1984), no. 3, 495–504.

[13] *English Title – Melodies online*, Retr. 27 August, 2007, from `http://www.dafos.dk/melodies-online.aspx`.

[14] *English Title – Melody codes and code searching*, Retr. 12 September, 2007, from `http://www.dafos.dk/melodies-online/melody-codes-and-code-searching.aspx`

[15] Eck, D., and N. Casagrande, "Finding Meter in Music Using An Autocorrelation Phase Matrix and Shannon Entropy", *Proceedings of 6th International Conference on Music Information Retrieval (ISMIR 2005)*, London, 2005, 504–509.

[16] Eerola, T., and P. Toiviainen, "MIR In Matlab: The MIDI Toolbox", *ISMIR 2004, 5th International Conference on Music Information Retrieval, Proceedings*, Barcelona, 2004.

[17] Elbourne, R.P., "The Question of Definition", *Yearbook of the International Folk Music Council* 7 (1975), 9–29.

[18] Elscheková, A., "Methods of Classification of Folk-Tunes", *Journal of the International Folk Music Council* XVIII (1966), 56–76.

[19] Erdély, S., "Classification of Hungarian Folksongs", *The Folklore and Folk Music Archivist* V (1962), 3, 1f.

[20] *Finnish Folk Tunes*, Retr. 1 March, 2007, from `http://www.jyu.fi/musica/sks/collection.html`.

[21] Garbers, J., *Integration von Bedien- und Programmiersprachen am Beispiel von OpenMusic, Humdrum und Rubato*, [Diss.], Fakultät IV – Elektrotechnik und Informatik der Technischen Universität Berlin, 2004 (`http://people.cs.uu.nl/garbers/publications/DissertationJoergGarbers.pdf`).

[22] Garbers, J., *et al.*, "On pitch and chord stability in folk song variation retrieval", [To appear in the Proceedings of the First International Conference of the Society for Mathematics and Computation in Music, Berlin], 2007.

[23] Hoos, H.H., K. Renz, and M. Görg, "GUIDO/MIR An Experimental Musical Information Retrieval System based on GUIDO Music Notation", *Proceedings of the 2nd International Symposium on Music Information Retrieval (ISMIR 2001)*, Bloomington (IN), 2001, 41–50.

[24] Huron, D., *Music Research Using Humdrum: A User's Guide*, Stanford CA, 1999. Online accessible at `http://csml.som.ohio-state.edu/Humdrum/guide.toc.html`. Retr. 27 August 2007.

[25] Huron, D., "The Melodic Arch in Western Folksongs", *Computing in Musicology* 10 (1995/6), 3–23.

[26] Jesser, B., *Interaktive Melodieanalyse*, Bern, 1991.

[27] Juhász, Z., "Contour Analysis of Hungarian Folk Music in a Multidimensional Metric-Space", *Journal of New Music Research* 29 (2000), nr. 1, 71–83.

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

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

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

[31] Keller, M.S., "The Problem of Classification in Folksong Research: a Short History", *Folklore* 95 (1984), no. 1, 100–104.

[32] Lartillot, O., "A Musical Pattern Discovery System Founded on a Modeling of Listening Strategies", *Computer Music Journal* 28 (2004), no. 3, 53–67.

[33] L.B. Meyer, *Style and Music — Theory, History, and Ideology*, Chicago, 1996.

[34] McCarty, W., *Humanities Computing*, Basingstoke, 2005.

[35] McNab, R.J., *et al.*, "The New Zealand Digital Library MELody inDEX", *D-Lib Magazine*, May 1997. Retr. 1 March, 2007 from `http://www.dlib.org/dlib/may97/meldex/05witten.html`.

[36] *[Without title]*, Retr. 1 March, 2007 from `http://www.nzdl.org/musiclib`.

[37] *The Mudcat Cafe*, Retr. 1 March, 2007, from `http://www.mudcat.org`.

[38] *Musipedia: The Open Music Encyclopedia*, Retr. 1 March, 2007, from `http://www.musipedia.org`.

[39] Nettl, B., *The Study of Ethnomusicology*, 2nd ed, Urbana, 2005.

[40] *Nederlandse Liederenbank*, Retr. 12 September, 2007, from `http://www.liederenbank.nl`.

[41] Peretz, I., and R. J. Zatore, "Brain Organization for Music Processing", *Annual Review of Psychology* 56 (2005), 89–114.

[42] Peretz, I., "The Nature of Music from a Biological Perspective", *Cognition* 100 (2006), 1–32.

[43] Pickens, J., "A Comparison of Language Modeling and Probabilistic Text Information Retrieval Approaches to Monophonic Music Retrieval", *ISMIR 2000, 1st International Symposium on Music Information Retrieval, Proceedings*, Plymouth, 2000.

[44] Rubin, D. C., *Memory in Oral Traditions*, New York, 1995.

[45] Schaffrath, H., *The Essen Folksong Collection*, D. Huron (editor), Stanford CA, 1995.

[46] Sapp, C., "Online Database of Scores in the Humdrum File Format", *Proceedings of 6th International Conference on Music Information Retrieval (ISMIR 2005)*, London, 2005, 664–665.

[47] Porter, J., "Convergence, Divergence, and Dialectic in Folksong Paradigms", *The Journal of American Folklore* 106 (1993), no. 419, 61–98.

[48] Singer, J., "Creating a nested melodic representation: competition and cooperation among bottom-up and top-down Gestalt principles", *ISMIR 2004, 5th International Conference on Music Information Retrieval, Proceedings*, Barcelona, 2004.

[49] Stein, E., "The Use of Computers in Folklore and Folk Music: a Preliminary Bibliography", *Library of Congress, Archive of Folk Song*, Washington, D.C., 1979. Online accessible at: `http://www.loc.gov/folklife/guides/BibComputers.html`. Retr. 27 August 2007.

[50] Steinbeck, W., *Struktur und Ähnlichkeit: Methoden automatisierter Melodieanalyse*, Kassel, 1982.

[51] Suchoff, B., "Editor's Preface", in: B. Bartók, *The Hungarian Folk Song*, Albany, 1981.

[52] Suppan, W., and W. Stief (eds.), *Melodietypen des Deutschen Volksgesanges*, Tutzing, 1976.

[53] Temperley, D., "A Probabilistic Model of Melody Perception", *ISMIR 2006, 7th International Conference on Music Information Retrieval, Proceedings*, Victoria (BC), 2006, 276–279.

[54] *Themefinder*, Retr. 1 March, 2007, from `http://www.themefinder.org`.

[55] Toiviainen, P., and T. Eerola, "Classification of Musical Metre with Autocorrelation and Discriminant Functions", *Proceedings of 6th International Conference on Music Information Retrieval (ISMIR 2005)*, London, 2005, 351–357.

[56] Typke, R., *Music Retrieval based on Melodic Similarity* [Diss.], Utrecht University, 2006.

[57] Volk, A., *et al.*, "Comparing Computational Approaches to Rhythmic and Melodic Similarity in Folk Song Research", [To appear in the Proceedings of the First International Conference of the Society for Mathematics and Computation in Music, Berlin], 2007.

[58] Wiora, W., "Systematik der musikalischen Erscheinungen des Umsingens", *Jahrbuch für Volksliedforschung* 7 (1941), 128–195.

[59] *YahMuugle*, Retr. 28 August 2007 from `http://pierement.zoo.cs.uu.nl/rism/index.html`.