TOWARDS INTEGRATION OF MIR AND FOLK SONG RESEARCH

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

Utrecht University and *Meertens Institute, Amsterdam

petervk@cs.uu.nl

ABSTRACT

Folk song research (FSR) 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 (MIR) research provides such methods. Yet a fruitful cooperation of both disciplines is difficult to achieve. We present a role-model to structure this cooperation in which tasks and responsibilities are distributed among the roles of MIR, Computational Musicology (CM) and FSR.

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 folksong recordings and transcriptions of the Meertens Institute (a research institute for Dutch language and culture in Amsterdam). Its purposes are on the one hand to support Folk Song Research (FSR) [1] 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. 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 form a so called 'tune family'.

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. As folk song melody research belongs to the domain of ethnomusicology, serious attempts to build software for processing folk song melodies should model concepts and methods that were developed in ethnomusicology. 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 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 describe an approach that may lead to better cooperation.

[1] The equivalent of the German 'Volksliedkunde'.

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

2 FOLK SONGS IN MIR

Only a limited number of MIR applications and studies are specifically aimed at searching folk song collections. Some online search engines allow the user to search in a large collection of folk song melodies. The *Danish Folklore Archives* [3] and the *Digital Archive of Finnish Folk Tunes* [6] are primarily meant for folk songs, while engines like *Themefinder* [19], *MELDEX* [12] and *Musipedia* [13] have a more general scope. Only the Danish search engine posesses 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.

Folk song melodies have been used as data in a considerable number of MIR studies. In some 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 [16]. In none of these papers the implications of the choice of this data set is discussed. In most cases it is simply stated that this 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.

In broader MIR circles some more studies have been done that particularly focus on folk songs. The work of Zoltán Juhász [10] is highly relevant in this respect. He selects his algorithms for their ability to answer questions about the data (mainly Hungarian folk song melodies) instead of employing the data to answer questions about his algorithms. By clustering contour representations of the melodies in various ways, his studies reveal differences between oral traditions in various countries.

Another relevant publication is an article by David Huron in which he proposes to visualize geographic differences in music by showing densities on a map [9]. As an example, using the Essen collection, Huron visualizes the geographic density of certain types of cadence notes, showing that Western European songs are nearly three times as likely to have most of their phrases end on a note other than the tonic compared to Eastern European songs.

Of the mentioned papers 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.

3.1 Oral Tradition

The transmission of songs in an oral tradition is determined by the capabilities of human perception, memory, performance and creativity. Participants in the 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 [15] shows that the representation of a song in human memory is not 'literal', but — in the words of Bertrand Bronson — "a fluid idea of a song" [2]. During performance the actual appearance of the song is recreated. In the process of 'translation' from the memory representation to audible words and melody, considerable variation may occur. As long as the the process of performing a song from memory is not sufficiently understood, we have mainly to focus on the recorded songs instances in order to understand this kind of variation.

A comprehensive inventory of types of variation in German folk songs is made by Walter Wiora [21], who summarizes the issue as: "Alles an der Beschaffenheit einer Melodie ist veränderlich".[3] He divides the types of change in seven categories: 1. changes in contour, 2. changes in tonality, 3. changes in rhythm, 4. inserting and deleting of parts, 5. changes of form, 6. changes in expression, 7. demolition of the melody. For each category he gives many examples.

3.2 Classification

A 'classification system' is used to group melodies with certain characteristics together. Examples are a common number of lines, a common number of syllables or a common cadence note sequence. Overviews of classification systems are provided by [4] and [1]. 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 by Ilmari Krohn. In his system the cadence notes (ending notes of the lines) are most important [5]. Béla Bartók and Zoltán Kodály adapted his system for Hungarian folk songs. In their publications songs were represented by: 1. the number of lines, 2. the sequence of cadence notes, 3. the number of syllables in each line, and 4. the ambitus [17, 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 [17, 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 [14, p. 124]. This applies to folk song classification systems in general [1, p. 33].

In the British-American folk song tradition, Bertrand Bronson found the following features to be important: 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. [2]. A classification system based upon these features can be expected to group songs in the same tune family together.

In [18] the German Archive of Folk Song (Deutsches Volksliedarchiv) uses an ordering based on the system of Krohn. The first criterion is the number of lines. Within the resulting groups the songs are ordered according to their cadence note sequences.

3.3 Identification

'Identification' of a song is related to the process of oral transmission. If two song instances are derived from the same common 'ancestor', they are considered to be the same song [14, p. 114].[4] 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. 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 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?

[3] Everything in a melody can change.

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

Which influences from abroad can be traced? How did the melodies develop over time?

At the Meertens Institute, the concept of 'melodienorm' (melody norm) [5] 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.

4 COOPERATION AND INTEGRATION

Despite some good examples presented in the previous sections, currently a profound mutual influence of MIR and FSR appears barely to exist. This seems to be true too for the relation between Musicology and MIR in general. Although the subject is the same (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 limited view on music prevents MIR often from being really 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 of MIR and musicology. The fundamental question is what we really want to achieve. Do we develop algorithms merely to *reproduce* a given 'ground truth', or do we 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 [11, 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, the most important concept to model is the melody norm.

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

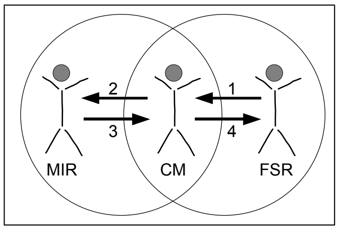

Figure 1. Three-role model for integration: Music Information Retrieval (MIR), Computational Musicology (CM), and Folk Song Research (FSR).

Although the modeling is more important than the models, for making the current state of knowledge available for application, these models are needed. This leads to two kinds of activities. First, the process of modeling and second, the implementation and deployment of the state-of-the-art models. These activities will alternate in an iterative process.

We now present a possible way to overcome the observed 'gap' with the help of the three-*role* model that is shown in Figure 1. 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 less commonly) the FSR and CM roles.

CM. In general the task for the CM-role is to connect the two disciplines. For the activity of modeling the task is to 'deconstruct' the 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 them to combine input, processing and output methods in various ways [7]. 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.

FSR. 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. 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. So, the task for FSR is to be as precise as possible in defining concepts. Another, more practical effort that may be expected from folk song researchers, is to take some time to learn how to use provided systems.

MIR. MIR can provide numerous useful software components and user interface components in which FSR concepts are implemented in an efficient and effective way. These components have to be packed into toolboxes or libraries by CM, thereby hiding implementation details that have no meaning in the musical domain and making components compatible with each other. In practice probably all MIR researchers play the role of CM to a certain extent. If they were not interested in music, they would not have been involved in MIR. The discipline of MIR can gain much from pursuing the CM-role more ambitiously.

5 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 ultimate 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 to just implement the models of e.g. Bartók or Bronson, since their feature sets were not created with the power of computational methods in mind, and they were often created for specific corpora. The possibilities Computer Science offers and the currently available computational power enable new kinds of models. Therefore, entirely different 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 [20, 8]. These new methods have to be developed in cooperation with musicologists who are able to provide the musical insights for modeling the features, and for improving failing models, thus escaping the problems of ground truths that were discussed in section 4. 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.

6 REFERENCES

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

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

[3] *Danish Folklore Archives — 10,000 melodies*, Retr. 1 March, 2007, from `http://www.dafos.dk/spmEnglish`.

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

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

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

[7] Garbers, J., *Integration von Bedien- und Programmiersprachen am Beispiel von OpenMusic, Humdrum und Rubato*, [Diss.], Berlin, 2004.

[8] 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], 2007.

[9] Huron, D., "Mapping European Folksong: Geographical Localization of Musical Features", *Computing in Musicology* 12 (1999–2000), 169–183.

[10] 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.

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

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

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

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

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

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

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

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

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

[20] 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], 2007.

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