

A PATTERN MINING APPROACH TO STUDY A COLLECTION OF DUTCH FOLK-SONGS

Peter van Kranenburg

Meertens Instituut
Amsterdam, Netherlands, and
Utrecht University
Utrecht, Netherlands

peter.van.kranenburg@meertens.knaw.nl

Darrell Conklin

University of the Basque Country UPV/EHU
San Sebastian, Spain, and
IKERBASQUE, Basque Foundation for Science
Bilbao, Spain

darrell.conklin@ehu.es

1. INTRODUCTION

In the ethnomusicological study of oral music cultures, the question what are the units of music has been of particular interest. Bohlman (1988) regards the song as the most basic unit. To better understand a given song culture, a possible next question would be what is the *smallest* unit of music. Nettl (2005, p.117) observes that folk musicians making field recordings are not always willing, or even unable to perform individual phrases, or motifs in isolation. Nevertheless, these units can to a certain extent have an independent existence, recurring in different pieces. This observation was first elaborated on by Tappert (1890), who entitled his study *Wandernde Melodien* (Wandering Melodies), employing the metaphor of traveling.

An important ethnomusicological concept we use in our study, is the concept of *tune family*, which has been introduced by Bayard (1950) to group together a set of folk song melodies that supposedly descend from one original tune through the process of oral or semi-oral transmission.

In a previous study on the way in which human collection specialists categorize Dutch folk song melodies into tune families (Volk & Van Kranenburg, 2012), it was found that the recurrence of short characteristic motifs is most relevant for the perception of similarity between songs belonging to the same tune family. Therefore, in the current work, we set out to analyse tune families in terms of shared melodic motifs.

In our approach, the set of melodies is divided into a *corpus* and an *anticorpus* (Conklin, 2010). The algorithm is capable of discovering recurring patterns that are statistically over-represented in the corpus with respect to the anticorpus. In all cases described in this paper, the corpus consists of all members of a given tune family, while the anticorpus consists of members of other tune families.

The question we ask is how to employ an existing sequential pattern mining algorithm (Conklin, 2010) to discover recurring patterns in a collection of Dutch folk tunes that can be considered building blocks for the melodies, and that characterize a melody as member of a tune family. In the following, we outline the method, the first results we obtained, and some open questions we want to address in our future work.

2. DATA

The pre-existing data set MTC-ANN 2.0, which is part of the Meertens Tune Collections (MTC) (Van Kranenburg et al., 2014)¹, contains 360 digitized vocal folk songs in 26 tune families from Dutch oral tradition, made available in symbolic encoding (*kern). These songs have been collected through ethnological field work in the Netherlands and from written sources such as song books. The collection specialists at the Meertens Instituut grouped the songs into tune families based on melodic similarity.

The small sample of 360 songs in 26 tune families has carefully been selected from a larger collection of thousands of songs. The sample is claimed to be representative for the larger collection concerning the kinds of variety that occur among variants of a tune family (Volk & Van Kranenburg, 2012). MTC-ANN 2.0 is provided with several sets of human annotations including a tune family label for each melody, but also 1,657 motif occurrences in 102 motif classes. Each of these motif classes represents an abstract melodic motif that has a number of concrete occurrences in songs within a tune family. These motifs are considered characteristic for the tune family in which they occur by the expert annotators. Therefore, we would expect an algorithmic pattern discovery method to find patterns that correspond to some of these annotated motifs.

3. METHOD

A melody is represented as a sequence of *events*, each a tuple comprised of basic attributes such as pitch, duration, and onset time. A *viewpoint* is a function that computes a value for each event in a sequence. Viewpoints can be *basic*: simply returning the basic attribute of an event; *derived* from other viewpoints; or *constructed*. For example, the derived level viewpoint, computed from the prevailing time signature and event onset time, describes the metric level of the event (0 being the highest metric level); and another derived viewpoint intref computes the diatonic interval from the reference pitch (the tonic) to the given pitch.

The choice of viewpoints is crucial for our study. The

¹ <http://www.liederenbank.nl/mtc>. Accessed: 5 June 2016.




Figure 1: Patterns discovered in tune family Koopman (top), and Stad (bottom), with one example occurrence. The colored notes constitute the occurrence, red indicating a note-event that is determined by non-pitch features only, green indicating the presence of pitch contour in the feature set, and blue indicating the presence of scale degree. The Koopman pattern describes a note that is the start of a phrase, followed by a note that is a major second above the tonic, and has shorter duration than the previous note, followed by a note of equal duration, and concluded with a note that has a higher pitch than the previous note. The Stad pattern describes a note somewhere in the middle of a phrase that is the fifth of the scale, and is approached by a leap of a third or fourth from the previous note, followed by, again, the fifth of the scale, then by the fourth of the scale, and concluded by a note of shorter duration, which is the first of a new phrase.

abstraction level of the viewpoints should be high enough to capture variability in the melodies as caused both by the process of oral transmission and by variations in choices that were made in the process of transcription into music notation. To achieve a suitable level of abstraction, we measure *relative* values for all viewpoints derived from pitch or duration.

For the current study we define the following viewpoints: *phrpos*, which records whether the note is the first in a phrase, the last in a phrase, or inside a phrase; *intref*, which represents the scale degree of the note given the key of the song; *c3i(level)*, which records whether the metric level of a note is higher, lower or equal with respect to the previous note; *c3(dur)*, which records whether the note is shorter, equal, or longer in duration than the previous note; *c3(pitch)*, which records whether the note is higher, equal, or lower in pitch than the previous note; *c5(pitch, 3)*, which records whether the note was approached by a leap (three semitones or larger), a step (smaller than a three semitones), or a unison, with distinction between ascending and descending intervals; and *c5(pitch, 7)*, which records whether the note was approached by a leap (seven semitones or larger), a step (smaller than seven semitones), or a unison, with distinction between ascending and descending intervals.

A *feature* is a tuple $\tau : v$ comprised of a viewpoint name τ paired with a value v . A *feature set* is a set of features, for example the feature set

$$\left\{ \begin{array}{l} c3(pitch) : - \\ intref : M2 \end{array} \right\}$$

contains two features, expressing that the pitch of the corresponding note is lower than that of the previous note, and is the major second (M2) of the scale. An event *instantiates* a feature set if all features in the set are true for the event.

A *feature set pattern* is a sequence of feature sets, and a song *instantiates* a pattern (or, stated equivalently, the pattern *occurs* in the song) if the successive feature sets of the pattern instantiate successive events in the song in at least one place. For example, the patterns shown in Figure 1

have four feature sets, with different features in each of them.

Following the method presented by Conklin (2010), a *one vs. all* strategy (Neubarth & Conklin, 2016) is used for mining patterns that contrast between groups of data. The method is designed to discover maximally general distinctive patterns (MGDPs), meaning that for each reported discovered pattern there is no more general pattern that is also distinctive. Each tune family is mined individually for distinctive sequential patterns, using each tune family F as a positive corpus and the rest of the pieces ($\neg F$) as the anticorpus.

In this work a statistical approach is used to measure the distinctiveness of a pattern: it is the probability p of finding at least the observed number of pieces of family F when taking a single random sample of pieces from the entire corpus $F \cup \neg F$. A pattern is then considered distinctive if its p-value falls below some specified significance level α (see Conklin, 2013, for details).

The MGDP set may contain overlapping patterns, so for the tune family mining task this set is further reduced by a greedy pruning strategy. Proceeding from the best (lowest p-value) pattern, a pattern is placed in the final set if it does not overlap, in any piece, with any pattern already in the final set. Thus none of the patterns in the final set will overlap in any piece with any other pattern.

4. RESULTS

The mining algorithm was applied repeatedly with each of the tune families in MTC-ANN in turn as corpus, while the other 25 tune families constitute the anticorpus. For this initial study, to obtain only a few highly distinctive patterns, we set the p-value threshold at the very low value of $\alpha = 10e-15$. The resulting set of discovered patterns contains 22 patterns in 14 tune families, showing that the algorithm is capable of discovering various kinds of melodic patterns that are significantly over-represented in the tune family.

We compare the discovered patterns with the manually annotated motifs as provided in MTC-ANN 2.0. These an-

notated motifs show what parts of the melodies are considered characteristic for the tune family according to human specialist annotators. We compute the establishment precision and recall² with a similarity function that considers an overlap of a discovered pattern occurrence with at least half of the notes of an annotated motif a hit, provided that the discovered pattern is not much longer than the annotated motif occurrence. We obtain an establishment precision of 0.86 and an establishment recall of 0.23, showing that the discovered patterns do correspond quite well with annotated motifs, but that the algorithm discovers much less patterns than human annotators did annotate. The low recall is caused by the very conservative p-value that we set. We only discover 22 patterns in 14 tune families, while the annotations consist of 102 motif classes in 26 families.

It is an open question what exactly this evaluation means. The motifs as provided in MTC-ANN 2.0 seem to be a highly subjective choice of the annotators. It is questionable to take this as *ground truth* for pattern discovery. Nevertheless, the high establishment precision suggests at least that the algorithm is able to find parts of the melodies that are considered stable within the tune family by human specialists. Further study of the interaction between the algorithmic results and the human annotations is needed.

As an example of a pattern that does not correspond with an annotated motif, we show a distinctive pattern that was discovered in tune family Koopman (adopting the abbreviations of the tune family names of Volk & Van Kranenburg, 2012). This pattern comprises an ascending contour starting from the tonic, which may seem trivial. However, the current results show that this particular way of starting a phrase is in fact rare outside Koopman.

The second example that is presented in Figure 1 is interesting because the fourth feature set of the pattern contains { phrpos : first }, which indicates a phrase break as part of the pattern. Such a phrase-crossing pattern would not be considered a motif in traditional hierarchical conceptualization of motifs in music theory. However, in the context of oral transmission, this seems a very meaningful piece of information, stating that this particular way of phrase transition, as part of the pattern, is specific for the tune family. For a singer generating a version of this tune, this might be crucial knowledge to properly sing the song.

5. CONCLUDING REMARKS

In this study, we present a first step towards a computational model of a given folk song culture as constituting of recombinations of a (possibly very large) number of melodic motifs. The occurrences of these motifs establish the identity of a song as member of a tune family. Since motifs may reoccur in a more or less varied appearance, the current approach in which not all notes of a motif are necessarily described with the same set of features, is very appropriate. The current study shows that the employed MGDG discovery method is capable of discovering parts

of the melodies that are stable within the variants of a tune family. Furthermore, it shows that the algorithmic mining results in patterns that very likely would not occur in traditional analysis, but that are meaningful in the context of understanding oral transmission of melodies.

There are several questions that should be addressed in future work when pursuing this approach. The relation of discovered patterns to experts' annotations is still poorly understood. Furthermore, there is still a gap between traditional musicological conceptualizations of motifs and tune families, and the kinds of patterns that are discovered by automatic discovery as presented in our study. We are convinced that a proper confrontation between the two domains will be beneficial for both, enriching traditional folk song analysis with objective methods, and enriching the algorithmic approach with knowledge of oral transmission of melodies.

6. ACKNOWLEDGMENTS

This research is partially supported by the project Lrn2Cre8 which is funded by the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET grant number 610859. Peter van Kranenburg is supported by the Computational Humanities Programme of the Royal Netherlands Academy of Arts and Sciences, under the auspices of the Tunes & Tales project. Thanks to Kerstin Neubarth for assistance with the manuscript.

7. REFERENCES

- Bayard, S. (1950). Prolegomena to a study of the principal melodic families of British-American folk song. *Journal of American Folklore*, 63(247), 1–44.
- Bohlman, P. (1988). *The Study of Folk Music in the Modern World*. Bloomington: Indiana University Press.
- Conklin, D. (2010). Discovery of distinctive patterns in music. *Intelligent Data Analysis*, 14(5), 547–554.
- Conklin, D. (2013). Antipattern discovery in folk tunes. *Journal of New Music Research*, 42(2), 161–169.
- Nettl, B. (2005). *The Study of Ethnomusicology: Thirty-one Issues and Concepts* (2nd ed.). Urbana and Chicago: University of Illinois Press.
- Neubarth, K. & Conklin, D. (2016). Contrast pattern mining in folk music analysis. In D. Meredith (Ed.), *Computational Music Analysis* (pp. 393–424). Springer.
- Tappert, W. (1890). *Wandernde Melodien: Eine musikalische Studie*. Leipzig: List und Francke.
- Van Kranenburg, P., De Bruin, M., Grijp, L. P., & Wiering, F. (2014). The Meertens Tune Collections. Meertens Online Reports 2014-1, Meertens Institute, Amsterdam.
- Volk, A. & Van Kranenburg, P. (2012). Melodic similarity among folk songs: An annotation study on similarity-based categorization in music. *Musicae Scientiae*, 16(3), 317–339.

² As defined at: http://www.music-ir.org/mirex/wiki/2015:Discovery_of_Repeated_Themes_&_Sections. Accessed: 5 June 2016.