Aspects of pattern discovery for Mozarabic chant realization

Darrell Conklin[1,2] and Geert Maessen[3]

[1] University of the Basque Country UPV/EHU, San Sebastian, Spain
[2] IKERBASQUE, Basque Foundation for Science, Bilbao, Spain
[3] Gregoriana Amsterdam, Amsterdam, The Netherlands

Abstract. Template methods for music generation can replicate some of the structure of an existing piece by making use of repeated patterns. New musical material sampled into repeated patterns of a template piece is repeated in the new piece. In this short paper we present some initial results on automating the process of intra-opus pattern detection in transcriptions of Mozarabic chant manuscripts.

Keywords: music generation, Mozarabic chant, pattern discovery

1 Introduction

From the sixth until the end of the eleventh centuries Christian worship on the Iberian peninsula was determined by the Mozarabic rite. The chant of this rite is called Old Hispanic chant. Only a few dozen Old Hispanic melodies are known with certainty. However, thousands of melodies of the old Mozarabic rite have been preserved in pitch unreadable, neumatic notation, the most important witnesses of which can be found in the early tenth century León antiphonary (Randel, 1973).

Though the pitches of the melodies are unknown and probably lost forever, the neumes provide important information to assist in their *realization*: determination of a singable and plausible pitch sequence representing the neumes. To achieve this, the neumes in a manuscript can be transcribed to contour notation. From note to note it is usually apparent if the melody goes up or down (Rojo & Prado, 1929). Figure 1 shows a fragment of the first part of the sacrificium *Aedificavit Moyses* for ordinary Sundays as copied in the León antiphonary (E-L 8, 306r). Shown at the top of the figure is one line of text with its neumes from the León antiphonary. Following that is the transcription of the neumes to contour letters (see Figure 1): `h`, a note higher than the previous note; `l`, lower; `e`, equal; and `o`, the start of a neume, with an unclear and undefined relative height to the previous note/neume. Within the contour sequence of Figure 1, neumes are separated by one dash, syllables by two dashes, and words by three dashes.

An open research task is the realization problem: creating plausible pitches for a specified neume contour sequence. In earlier work (Conklin & Maessen, 2019) a general method for generating such melodies was described. The basis

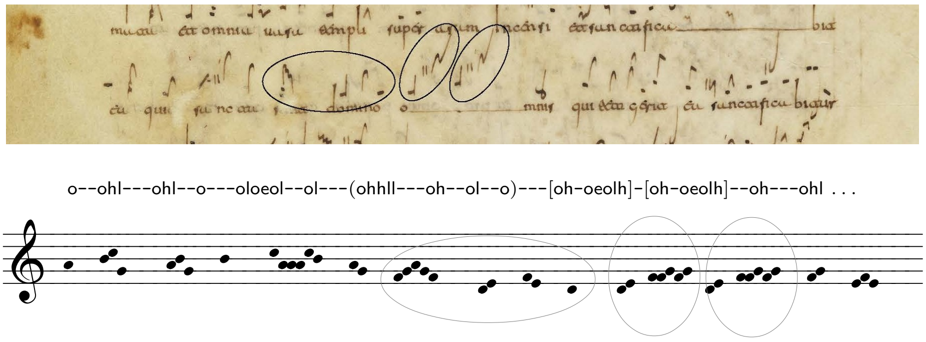

Fig. 1. The first part of the sacrificium *Aedificavit Moyses* for ordinary Sundays as copied in the León antiphonary (E-L 8, 306r). In the neumated text (*ea quia sancta sunt domino omnis qui tetigerit ea sanctificabitur*) three instances of two patterns are encircled (referred to in Section 3 as C, A, and A). Following this is the contour transcription for the first half, with instances of the patterns bracketed. Bottom: a generated realization of the contour sequence into pitches.

of this is to represent the contour sequence of a melody as a *template* to which generated pitch sequences must conform. Templates are filled by Monte Carlo sampling from a generative statistical model trained on a corpus of pitched chants. Templates additionally reference regions of intra-opus repetition which give coherence to a melody, either on long repeated sections, or on much shorter regions where repetition is more subtle. The central assumption made is that all instances of the same intra-opus pattern will be instantiated by the same pitches in any performance.

In this extension to earlier work on chant realization (Conklin & Maessen, 2019; Maessen & van Kranenburg, 2017), an automated pattern discovery method is used to find intra-opus patterns. A solution could free the musicologist from the intricate task of finding patterns in long neume transcriptions. Additionally, potentially more patterns can be found, including ones that are subtle and perhaps less obvious by visual inspection alone.

This short paper will report work in progress on the integration of pattern discovery into chant realization using templates. The results might also be relevant to any template-based method for music generation (Collins et al., 2016; Padilla & Conklin, 2018; Cope, 2001), as they indicate how classical concepts of sequential pattern mining might be insufficient, and how careful attention to structural boundaries and statistical significance are necessary.

2 Pattern discovery

Mozarabic chant realization can be performed using a template method (Conklin & Maessen, 2019). Templates are sequences of features, each feature specifying the desired properties of a particular music event. Features include contour relations, and note ranges (specified manually), defined (fixed) pitches, and indeed any viewpoint of an event (Conklin & Witten, 1995). Features may reference *variables*, which specify shared content with other events.

2.1 Patterns

In the most general case, a pattern could be any sequence of contour letters. It was found, however, that this definition is too general because it allows pattern instances to end within neumes. Therefore for the purpose of this study, all instances of a pattern should begin and end on neume boundaries (indicated by o, see Figure 1). For example referring to the fragment in Figure 1, ohhlloholo spans four neumes. On the other hand, neither prefix ohhllo nor ohhl can be viable pattern instances as they do not end at a neume boundary.

An *instance* of a pattern Φ in a sequence is a contiguous subsequence that matches the pattern. Note that overlapping instances are accommodated. The *total count* $c(\Phi)$ of a pattern Φ is its total number of instances: the number of times the particular neume sequence occurs in the template. A pattern Φ is *contained* in a pattern Φ' if $\Phi' = X\ \Phi\ Y$ for (possibly empty) patterns X and Y.

2.2 Significant closed patterns

A huge number of potential patterns can exist in a given sequence and it is necessary to restrict this space, both structurally (by the types of patterns found), and statistically (as a filter on the pattern set).

Two common structural restrictions are *closed* and *generator* patterns (Fournier-Viger et al., 2017). Generators cannot be pruned to a pattern with the same count, whereas closed patterns cannot be extended without a drop in the total count. They therefore represent the top, and the bottom, of containment chains of patterns with the same total count. For inter-opus pattern discovery, usually the shortest, most general significant patterns are desired (Conklin, 2020). For intra-opus pattern discovery, as noted by other researchers (Lartillot, 2005), closed patterns are more appropriate as they cover the most musical material possible without becoming overly specialized.

More precisely, a *closed pattern* Φ is one which is not contained in a longer pattern Φ' with $c(\Phi') = c(\Phi)$. Thus the closed patterns are the longest patterns among containment chains of patterns with the same count. To find all closed patterns in the neume contour sequence, we implemented an algorithm based on the bidirectional closure checking algorithm (Wang, Han, & Li, 2007), adapted here for intra-opus patterns. The algorithm performs dynamic pruning of the search space to efficiently omit entire paths that cannot possibly lead to closed patterns.

The set of closed patterns, while smaller than the set of all possible patterns, may still be very large and it is necessary to prune these to patterns that are significant: those patterns unlikely to occur with the observed count by chance alone (Goienetxea & Conklin, 2018). To do this, an expected total count λ is

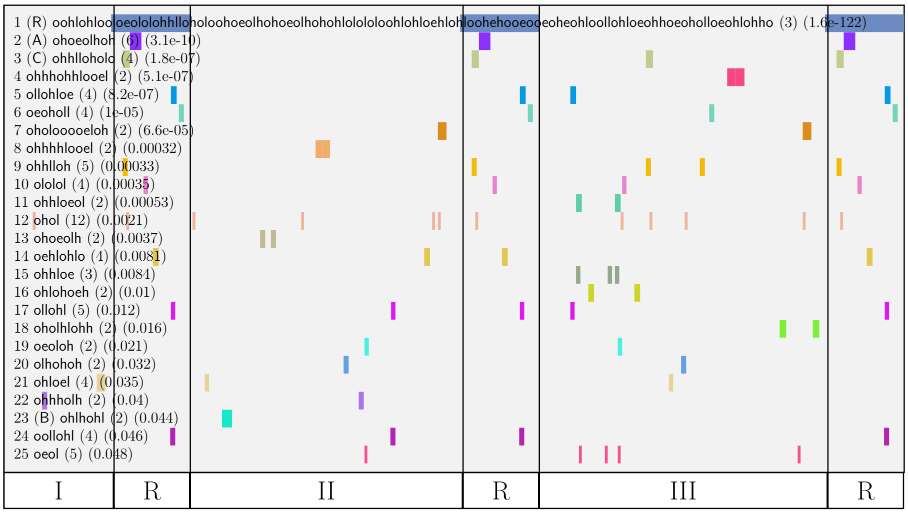

Fig. 2. Patterns discovered at $\alpha = 0.05$. Part boundaries are indicated at the base.

computed using a zero-order background model of neumes derived from the analysis piece. The right tail of the Poisson distribution with parameter λ is applied to the observed count to compute a p-value.

Following pattern discovery, a template is created, ensuring that all instances of the same pattern are represented by the same sequence of variables.

3 Results and discussion

The chant *Aedificavit Moyses* was used as a case study. This has the form I-R-II-R-III-R, where I is the initium, R is the repetendum, and II and III are the second and third parts. It has a total of 1251 notes within 471 neumes. With a minimum total count of 2 and no significance pruning, 890 patterns can be found of which 125 are closed. Filtering this with a p-value upper threshold of $\alpha = 0.05$ produces 25 closed significant patterns. Figure 2 shows the set of discovered patterns, sorted by increasing p-value. Some are very repetitive over the entire piece. 461 notes (or 0.37 of the entire piece) are covered by one or more of the 25 closed significant patterns.

While recognizing the possibility of additional patterns revealed by pattern discovery, at least the following four annotated patterns should be revealed by a pattern discovery algorithm:

- pattern R: the *repetendum*: a long pattern of over 100 notes (3 instances);
- pattern A: ohoeolh (6 instances);
- pattern B: ohlhohl (2 instances);
- pattern C: ohhlloholo (4 instances).

As expected, the long repetendum R appears at rank 1. One minor change to the annotated repetendum is the presence of two neumes (ohl and o) at the left hand part of the pattern. Since in the manuscript the repetendum is only written and notated in its first instance, this is significant. In fact these two neumes (four notes) give the audible sign to the singers that the repetendum should be repeated. Thus the pattern discovery, not informed by segment boundaries, has produced a slightly longer pattern. The right hand side of the repetendum is exact. Note that applying *covering* — building a pattern set where only one pattern covers any particular region of the piece (Padilla & Conklin, 2018) — is not applicable as it will filter the nested repetendum patterns.

Pattern C, at rank 3, appears exactly as specified. It occurs in all three repetendi and also once in verse III, providing coherence between the two parts of the chant. Pattern A appears on rank 2 as ohoeolhoh. The two circled instances in Figure 1 correspond to the first two instances in Figure 2. It is the desired A pattern extended by one neume (oh). On close inspection of this pattern it was found that this extension by one neume causes the pattern to span two types of boundaries. It can be seen in Figure 1 that its first instance spans three neumes within one syllable, where its second spans three neumes on two syllables. This example indicates that patterns should ideally respect the structural levels of a chant: syllable, word, sentence, and part. For example, referring to pattern A, the trailing oh would be omitted as it induces the second instance of the pattern to span a syllable boundary. Future work could investigate an extension of the pattern discovery algorithm to handle annotated level boundaries.

The issue with pattern A could be solved by considering generator rather than closed patterns, i.e. ohoeolh is a generator pattern. However this is not generally an acceptable solution as generator patterns would not, for example, cover the entire repetendum (there are many non-closed subpatterns of the repetendum that also have a total count of 3). Note that the discovered version of pattern A interestingly has self-overlap (ending and beginning on the same neume oh). In purely descriptive mining (patterns inspected by an analyst), self-overlapping patterns should be avoided as they are hard to interpret. Future work will examine whether generation will in fact benefit from including self-overlapping patterns or whether the template might be slightly over-constrained.

Pattern B does not appear in the top results but much lower at rank 23. It occurs twice, in a contiguous clump, in verse II. The method for p-value computation may need to be adjusted, as one would hope that the B pattern would appear with higher rank. For future work this could involve adjusting the expected count definition.

Regarding all other patterns discovered, many of these may be significant, even the short ones spanning just a few neumes (e.g., patterns 12, 25). Future work will explore whether these short patterns over-specify the template for chant realization.

Acknowledgements Thanks to Kerstin Neubarth for valuable discussions on the research and the manuscript.

References

Collins, T., Laney, R., Willis, A., & Garthwaite, P. H. (2016). Developing and evaluating computational models of musical style. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, *30*(1), 16–43.

Conklin, D. (2020). Patterns and antipatterns for music corpus analysis. In D. Shanahan, A. Burgoyne, & I. Quinn (Eds.), *Oxford Handbook of Music and Corpus Studies.* New York: Oxford University Press. (in press)

Conklin, D., & Maessen, G. (2019). Generation of melodies for the lost chant of the Mozarabic rite. *Applied Sciences*, *9*(20).

Conklin, D., & Witten, I. (1995). Multiple viewpoint systems for music prediction. *Journal of New Music Research*, *24*(1), 51-73.

Cope, D. (2001). *Virtual Music: Computer Synthesis of Musical Style.* Cambridge, MA: The MIT Press.

Fournier-Viger, P., Lin, J., Kiran, R., Koh, Y., & Thomas, R. (2017). A survey of sequential pattern mining. *Data Science and Pattern Recognition*, *1*(1), 54-77.

Goienetxea, I., & Conklin, D. (2018). Melody transformation with semiotic patterns. In M. Aramaki, M. E. P. Davies, R. Kronland-Martinet, & S. Ystad (Eds.), *Music technology with swing* (pp. 477–488). Cham: Springer International Publishing.

Lartillot, O. (2005). Efficient extraction of closed motivic patterns in multi-dimensional symbolic representations of music. In *Proceedings of the 6th International Conference on Music Information Retrieval (ISMIR 2005)* (pp. 191–198). London, UK.

Maessen, G., & van Kranenburg, P. (2017). A semi-automatic method to produce singable melodies for the lost chant of the Mozarabic rite. In I. Barbancho, L. Tardón, & A. Peinado (Eds.), *Proceedings of the 7th International Workshop on Folk Music Analysis* (pp. 60–65). Malaga.

Padilla, V., & Conklin, D. (2018). Generation of two-voice imitative counterpoint from statistical models. *International Journal of Interactive Multimedia and Artificial Intelligence*, *5*(3), 22-32.

Randel, D. (1973). *An Index to the Chant of the Mozarabic Rite.* New Jersey: Princeton University Press.

Rojo, C., & Prado, G. (1929). *El Canto Mozárabe, Estudio histórico-critico de su antigüedad y estado actual.* Barcelona: Diputación Provincial de Barcelona.

Wang, J., Han, J., & Li, C. (2007). Frequent closed sequence mining without candidate maintenance. *IEEE Transactions on Knowledge and Data Engineering*, *19*, 1042-1056.