TWO METHODS TO COMPUTE MELODIES FOR THE LOST CHANT OF THE MOZARABIC RITE

Geert Maessen
Gregoriana Amsterdam, Amsterdam, The Netherlands
gmaessen@xs4all.nl

Darrell Conklin
Department of Computer Science and Artificial Intelligence, University of the Basque Country UPV/EHU, San Sebastian, Spain
IKERBASQUE, Basque Foundation for Science Bilbao, Spain
darrell.conklin@ehu.eus

ABSTRACT

Several medieval chant traditions are preserved in precursors of modern music notation. Virtually all chants of the Mozarabic rite are only preserved in the earliest of these: pitch-unreadable neumatic notation. Melodic intervals are not available. This paper sketches two computational methods to produce melodies based on a comparison of transcriptions of the early notation with pitch-readable preserved traditions encoded in a data set.

1. NOTATION, ENCODING & AN EXAMPLE

In 1973 Don Randel published a detailed description of over 5,000 chants of the Mozarabic rite preserved in about forty manuscripts and fragments dating from the early eighth until the thirteenth centuries. Several genres were included, from very simple to most complex: Randel orders 28 genres in five manuscript groups with neumatic notations: León, Rioja, Silos, Toledo A and Toledo B. The most important manuscript is the León antiphoner (E-L 8) dating from the early tenth-century. Unlike Gregorian chants, many Mozarabic chants appear in only one or two manuscripts, often with much greater differences in musical detail than in Gregorian chant. Only a few dozen relatively simple melodies have been found in pitch-readable notation. Some scholars, however, have shown melodic relations with other chant traditions for some specific Mozarabic chants (Levy, 1998). Therefore we have compiled a data set with preserved medieval melodies, as a base for the construction of melodies for the lost Mozarabic chant (Van Kranenburg & Maessen, 2017).

Although melodic information is virtually absent, the neumatic notation of the Mozarabic rite sketches the contours of the melodies. From note to note we can mostly see if the melody goes up or down (Rojo & Prado, 1929). An elementary way to represent this contour information is using six letters: `h` a note higher than the previous note; `l` a note lower; `e` a note of equal pitch; `b` higher or equal; `p` lower or equal; `o` a note with unclear relative height. Figure 1 shows the beginning of the second part of one of Levy's chants; the sacrificium *Sanctificavit Moyses altare*. Shown at the top of the figure are three lines from the León antiphoner. Following that, three parallel lines show the transcription of the neumes to contour letters and two different melodies produced with our two methods for chant reconstruction. Encircled in the manuscript image is a repeated *intra-opus* neumatic pattern that should be instantiated by the same musical material. In the contour string and melodies the corresponding patterns are underlined. Capitals in the string indicate notes of the patterns.

It should be noted that the context of this particular pattern is interesting for the comparison of chant traditions. The chant text is a narrative from Exodus (Ex. 34:2-5): The Lord said to Mozes "come up to me unto mount Sinai" (ascende ad me in montem Syna). Then Mozes went up unto the mount (ascendit in montem) and the Lord descended towards him (descendit ad eum). The three words "ascende", "ascendit" and "descendit", share the pattern. This pattern is also part of two extended patterns; the first shared by "ascende" and "ascendit", the second by "ascendit" and "descendit". So here the music can be seen to express the "meeting" of the Lord and Mozes, something hardly imaginable in other chant traditions. A closer look may even reveal a metaphor for the idea that Mozes is in God's "hands": "His" words and deeds, "ascende", and "descendit". Misunderstanding of this kind of sophistication may well have been used in the repression of the Mozarabic rite in the late eleventh century when, in the ongoing power struggle between the advocates of the different rites, the supposed heretical character of the Mozarabic rite was repeatedly stressed by Pope Gregory VII (Vones, 2007). In the parallel chants of the surviving traditions these details are blurred, completely absent, or at best reduced to different up and down movements only.

2. METHOD 1

The first method to produce melodies for the lost chant has been described in detail in a previous paper (Maessen & Van Kranenburg, 2017) and is reviewed here briefly. In search for a melody of a specific lost chant, we first transcribe the chant notation to a string of contour letters. Then we divide this string into segments, guided by the grammatical structure of the chant. Our method implemented a brute force string matching algorithm that searches all matches of all segments of the contour string in all chants of the data set, allowing for a variable number of n skips in all segments. The algorithm then lists the best matching melodies of the data set conforming to a computed evaluation score S. This score is defined by the positions of the matches of the segments in the data set melodies and the number n of allowed skips.

In order to produce a singable melody for the lost chant we may need three additional steps using this method. The first combines the matches of the segments in the best data set melody to one melody for the lost chant. For some of the segments we may need to increase the number of allowed skips n, or shorten the segments. In a second step we may give repeating patterns (*intra-opus* patterns, as appearing in the early notation of a single chant) the same pitch sequences. The third step consists in rehearsing the chant and correcting some of its "uncharacteristic" pitches. Some segments may be in need for transposition, and especially at the borders of segments melodic lines sometimes need to be smoothed.

With this method we produced over 100 chants, some of them still on the internet (Gregoriana Amsterdam, n.d.). Although some of these chants are considered beautiful, there are some problems (see Section 4).

3. METHOD 2

The basis for the second method is the construction of statistical models of a coherent corpus, and then "inverting" this model to generate new music having high probability according to the model (Conklin, 2003). A statistical model is trained on a data set of 137 Gregorian offertories, comprising a total of approximately 65,000 notes. Given the size of the corpus, it was possible to create a bigram model of pitches that does not have data sparsity problems. Following model training, a sequence of pitches can be generated based on the probabilities derived from the data set by performing statistical Gibbs sampling and settling on sequences at the high end of probability space.

To capture the specified positional constraints and also *intra-opus* repetition, Conklin (2016, 2017) introduced an important improvement to statistical generation that makes it appropriate for the production of melodies for the lost chant at hand. Template pieces are encoded using a *semiotic pattern*, which specifies constraints on individual notes and also equality relations between segments of notes. Sequences are sampled from the trained statistical model while ensuring that the semiotic pattern is maintained for each sample. Positional constraints, the most important facet of the pattern in the case of chant reconstruction, are given by the string of contour letters referred to in Section 2 above and illustrated in Figure 1. Repeating (*intra-opus*) contour patterns are generated as patterns of equal pitches, for example, see the three (overlapping) repeated patterns in capitals in Figure 1. These *intra-opus* patterns are annotated manually for each template.

Given a model and a semiotic pattern, a large space of sequences is sampled in two steps. In the first, an iterative random walk (Conklin, 2016) is used to create an initial solution compatible with the semiotic pattern. In this step the contour pattern, using the specified ambitus desired, is compiled into further positional constraints that make the search for an initial solution feasible. For example, if an `h` contour is at a position, it is clear that the previous position cannot be at the maximum height of the ambitus.

Following the successful production of an initial solution, Gibbs sampling is performed: positions are selected uniformly over the semiotic pattern variables, and all new possible notes are considered in that position. A new sequence is then sampled from the distribution of pieces with successfully substituted notes. Unlike standard Gibbs sampling, here some preference is given to sampling sequences that increase rather than decrease the current sequence probability. In our implementation, we have found that, for most templates, after approximately 100,000 iterations the highest probability solutions no longer change.

Until now we successfully generated several chants using positional and *intra-opus* patterns together. Some of these melodies we performed in videos with the early notation running along (Gregoriana Amsterdam, n.d.). We also successfully experimented with other constraints, such as the ambitus (different for different parts), and the first and last (and other) pitches of the chant.

4. COMPARISON

As Figure 1 and Gregoriana Amsterdam (n.d.) may show, both methods are able to generate singable melodies. Of importance here, however, are the differences:

1. Due to the segmentation and the working hypothesis of the existence of (nearly) identical matches of segments to parts of existing melodies, Method 1 seems infeasible without manual editing of the constructed melody. In Method 2 manual editing is not necessary because, unlike Method 1, there are no "problematic" borders between segments. Any sequence following the semiotic pattern will be a solution to the contour and *intra-opus* pattern specification.

2. The score S of Method 1 provides a good criterion for the relation with the lost melody. When S is higher than about 70 %, Method 1 also indicates serious candidates for historically related melodies in the data set. However, until today, Method 1 seldom produced a score above 40 %. Method 2 is only able to generate general characteristics about probabilistic space as it uses only a low order statistical model.

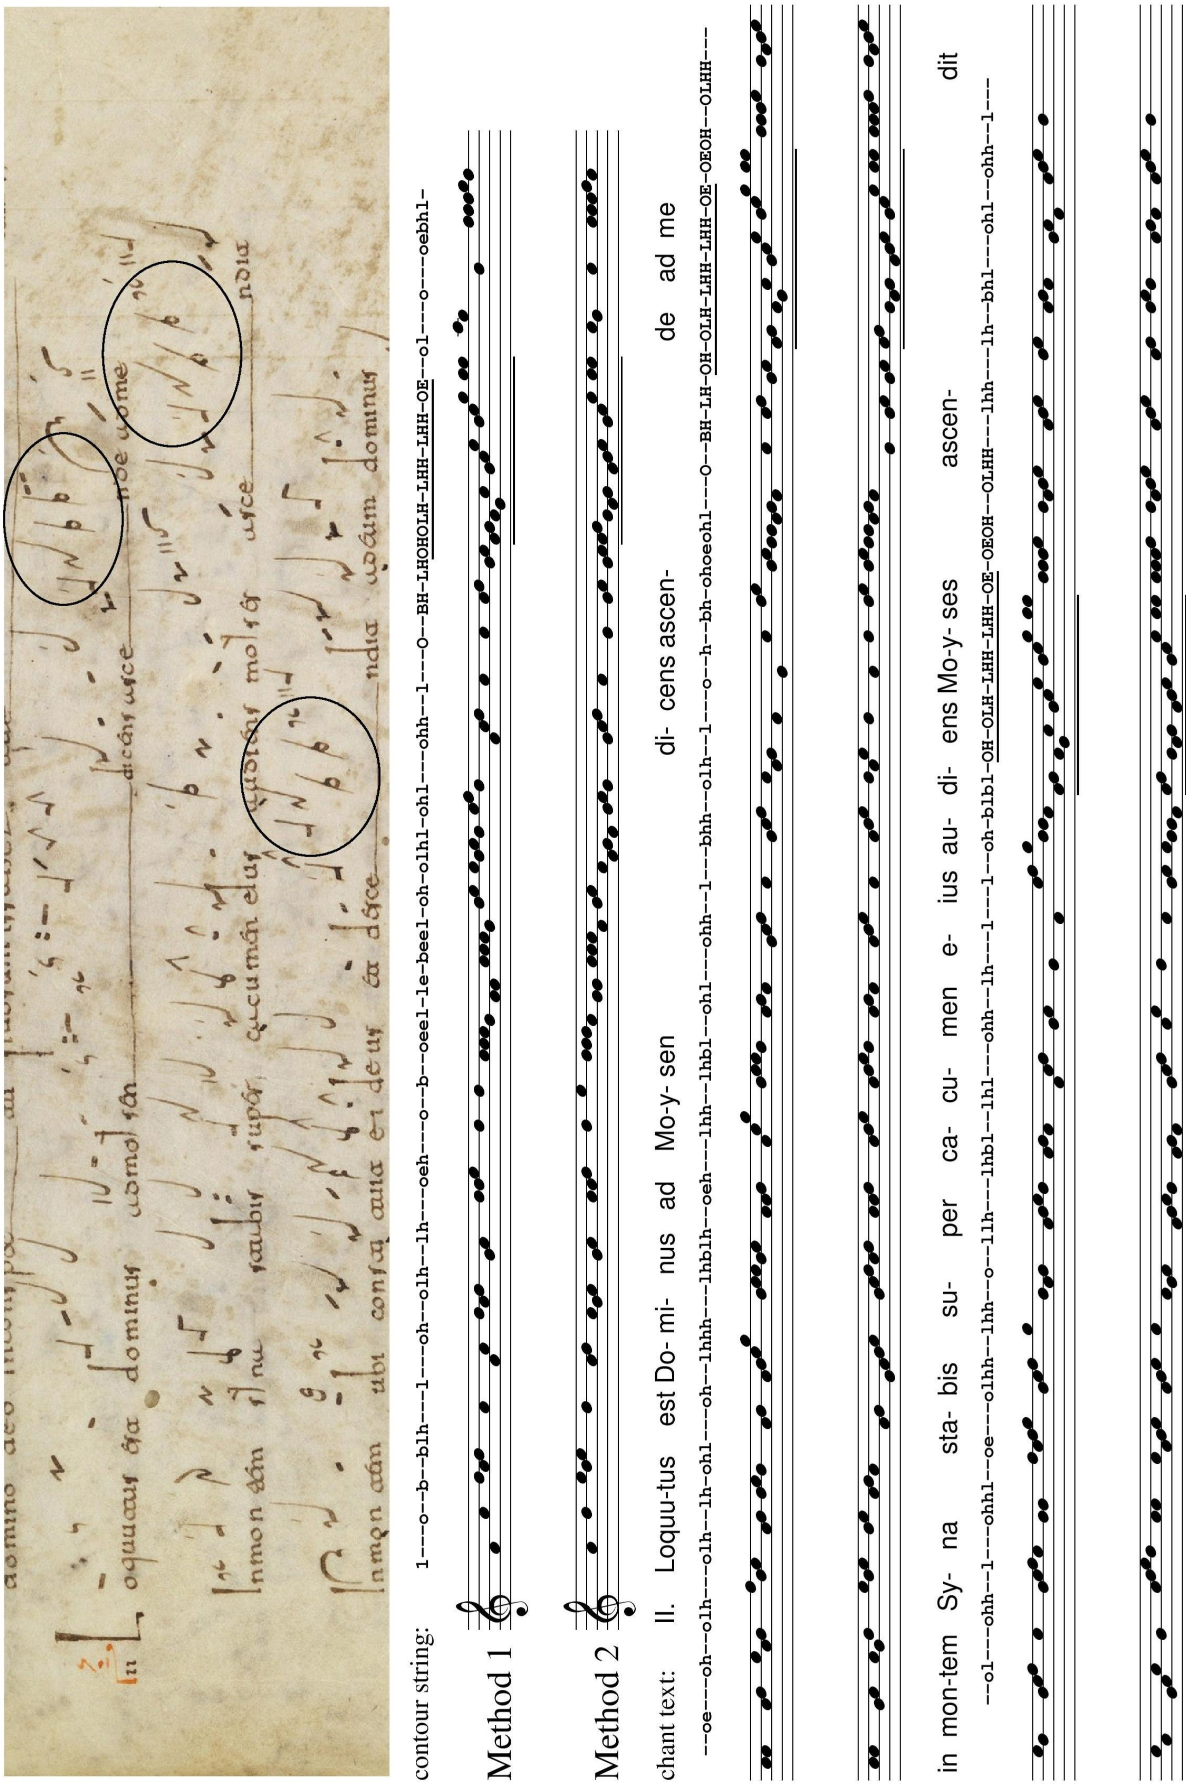

Figure 1. A passage from *Sanctificavit Moyses:* neumatic notation (E-L 8; 305r13), contour string and two melodies.

3. In Method 1 we did not yet fully use the information included in the contour string about syllables, words and sentences. In Method 2 this information is automatically processed as part of the statistical model.

4. The more constraints we introduce to Method 1 (patterns, ambitus, specific pitches), the more problematic it will be to construct a melody, unless the desired melody (or a very close variant) is included in the data set. In Method 2 the only things of importance are the sampling algorithm and the statistical model used.

5. In order to construct a considerable self-consistent corpus of chants it will be necessary to produce *inter-opus* patterns of equal pitches, i.e. to generate patterns occurring in different chants and having equal sequences of pitches. Given its reliance on segmentation and the facts relating to its general hypothesis this seems almost impossible in Method 1. However, we are already successfully experimenting with this in Method 2.

5. CONCLUSION AND FUTURE WORK

The overall impression is that both methods are able to generate singable melodies. The second, however, even in this stage, seems less laborious. No manual editing is needed and the generation of *intra-opus* repeating patterns is implicit in the method.

A fascinating point opened up by our research is the role of overfitting in statistical models. Usually this is viewed negatively as the inability of a model to generalize past the known data. However in the chant reconstruction problem there are cases where overfitting is desired, as for example when a template melody may contain *inter-opus* patterns (seen in another chant). These patterns should be used when available. Thus high on our agenda is the consideration of how to handle *inter-opus* recurring patterns. In summary Method 2 seems not only the best option to generate unique chants, but also to construct a self-consistent repertory agreeing with the early notation, something that seems hardly feasible with the first method. And, of course, this last option is high on our agenda.

Presently we are, therefore, working on the implementation of *inter-opus* patterns. We are also constructing some cross-validation cases: chants where neumes and corresponding pitches are known, e.g. the Gregorian chant offertory *Scapulis suis* (Gregoriana Amsterdam, n.d.). There are other items on our agenda. Pattern discovery algorithms might be used to find *intra-opus* patterns in the templates, thus automating the laborious step of hand annotation of a template for patterns. To create large collections of reconstructions for many templates this seems even necessary. In Method 2 it will be necessary to handle church modes and ambitus constraints: these may vary throughout the piece and will require some broad segmentation of the template. Also in Method 2 the reference of patterns to the conservation of exact pitch sequences seems unnecessarily strict and we plan to allow any feature (intervals, neume shapes) to be conserved between pattern instances. Until now Method 2 only made use of a single tradition. Since we know that several traditions were related to the lost chant (Levy, 1998), it will be necessary to handle the differences between these traditions in Method 2. We are working on ways to define the relations between the lost chant and these traditions. Finally, until now we only used six contour letters. However, the neumatic information in the León antiphoner is much richer. Since the meaning of Mozarabic neumes is similar to Gregorian neumes and we do know the Gregorian melodies, it will be wise to include still another data set in our algorithms: Gregorian chants in pitch-unreadable neumes.

6. REFERENCES

Conklin, D. (2003). Music generation from statistical models. In *Proceedings of the AISB Symposium on Artificial Intelligence and Creativity in the Arts and Sciences, Aberystwyth, Wales,* (pp. 30-35).

Conklin, D. (2016). Chord sequence generation with semiotic patterns. *Journal of Mathematics and Music*, 10(2):92-106.

Conklin, D. (2017). Modelling and sampling jazz chord sequences. In *Proceedings of the 10th International Workshop on Machine Learning and Music, 6.10.2017, Barcelona, Spain,* (pp. 13-18).

Gregoriana Amsterdam. (n.d.). YouTube Channel of Gregoriana Amsterdam, with some sung examples produced with both methods (video titles with numerals 1 and 2) and showing the original notation: https://youtube.com/lelalilu

Levy, K. (1998). Toledo, Rome and the Legacy of Gaul. In *Gregorian Chant and the Carolingians*, 31-81.

Maessen, G. & Van Kranenburg, P. (2017). A Semi-Automatic Method to Produce Melodies for the Lost Chant of the Mozarabic Rite. In *Proceedings of the 7th International Workshop on Folk Music Analysis, 14-16 June 2017, Málaga, Spain,* (pp. 60-65).

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.

Van Kranenburg, P. & Maessen, G. (2017). Comparing Offertory Melodies of Five Medieval Christian Chant Traditions. In *Proceedings of the 18th International Society for Music Information Retrieval Conference, Suzhou, China, October 23-27,* (pp. 204-210).

Vones, L. (2007). The Substitution of the Hispanic Liturgy by the Roman Rite in the Kingdoms of the Iberian Peninsula. In *Hispania Vetus: Musical-Liturgical Manuscripts: from Visigothic Origins to the Franco-Roman transition,* 43-59.