Discovery of generalized interval patterns

Darrell Conklin and Mathieu Bergeron
Music Informatics research group
Department of Computing
City University, London
United Kingdom

***Abstract*— This paper presents a method for pattern discovery based on viewpoints and feature set patterns. The representation for pattern components accommodates in a fully general way the taxonomic relationships that may exist between interval classes. A heuristic probabilistic hill climbing algorithm is developed to rapidly direct the search towards interesting patterns. The method can be used for single piece analysis, for comparison of two pieces, and also for pattern analysis of a large corpus. The method is applied to the music of French singer-songwriter Georges Brassens.**

***Keywords*— pattern discovery, melodic interval, viewpoints, feature set, subsumption.**

I. Introduction

Pattern discovery is an important part of computational music analysis systems. Patterns repeated in a single piece can provide information about its structure and motivic development, while those found in two pieces can indicate a deeper similarity in musical material. Patterns found to be recurrent in a large corpus may indicate general aspects of the genre or style under analysis.

Pattern discovery methods rely centrally on the representation used for pattern components. Common representations for melody include various types of melodic interval and contour ([4], [5], [8]). For example, the standard three-point melodic contour ***c3*** partitions all melodic intervals into three disjoint classes $\{d, r, u\}$ representing a motion down, a repetition, or a motion up; the more specialized five-point contour ***c5*** into five classes $\{-l, -s, r, +s, +l\}$ representing leap down, step down, repetition, step up, and leap up. Steps involve a motion of 1 or 2 semitones, and leaps a motion greater than 2 semitones.

In computational music analysis it is often difficult to choose *a priori* the correct level of abstraction needed for pattern components. For example, melodic interval ***int*** may be too specific to reveal patterns, while the melodic contour ***c3*** may not have sufficient specificity. However, the choice of any single representation can lead to a loss of generality or specificity (for example, with the ***c5*** representation, the more general ***c3*** classes and the specific intervals represented by ***int*** are not accessible).

Figure 1 illustrates this problem with a musical example; the first two phrases of the French singer-songwriter Georges Brassens' *Le Bistrot* [3], encoded using ***int***, ***c5***, and ***c3***. The apparent melodic similarity between the two phrases cannot be captured by any single viewpoint of analysis. For example, the intervals -2 and -1 (second last note) do not match using the ***int*** representation, but do match using ***c5*** (they are both a "step down"). The intervals of $+3$ and $+1$ (note 11) do not match with ***c5*** but do match with ***c3*** (they are both upwards melodic motions). Some intervals require abstraction further to that captured by ***c3*** and ***c5***. For example, the intervals of $+2$ and -2 (note 5) are both "steps", and $+3$ and -1 (note 6) are both "any interval".

The focus of this study is on *generalized intervals*, where partitions of melodic intervals are overlapping and have subset relationships that can be represented in a taxonomy. Generalized intervals can be represented as *feature sets* (Figure 2), where each feature set is a logical conjunction of interval classes. In this scheme, feature set specialization (subsumption) is a subset relation. For example, $\{\boldsymbol{s}\}$ (step) subsumes $\{\boldsymbol{s}, \boldsymbol{u}\}$ (step up) and $\{\boldsymbol{s}, \boldsymbol{d}\}$ (step down). Note that the taxonomy in Figure 2 is not a full lattice, as some feature sets (e.g., $\{\boldsymbol{s}, \boldsymbol{r}\}$) cannot be satisfied by any interval and are therefore contradictory.

For brevity, in Figure 2 the taxonomy is not presented to the depth of melodic intervals; for example, it is understood that the feature set $\{\boldsymbol{l}, \boldsymbol{u}\}$ (leap up) subsumes all feature sets representing melodic intervals greater than $+2$ (e.g., the set $\{\boldsymbol{l}, \boldsymbol{u}, +3\}$). The feature set representation elegantly captures all of the ***c3*** and ***c5*** classes, in addition to the ***s*** (step) and ***l*** (leap) interval classes.

This paper presents an efficient method for the discovery of patterns where pattern components are feature sets and there are subsumption relations between the sets. The remainder of this paper describes the representation in more detail, and develops and applies a pattern discovery algorithm that searches for maximal patterns with generalized intervals as components.

II. Methods

The method for pattern representation and discovery is based on the concept of the *feature set pattern* as reported earlier [1]. There, the notion of a *maximal* pattern, which is a frequent pattern that cannot be further specialized, was used within a pattern discovery algorithm. Maximal feature set pattern discovery can be computationally expensive, because there are many subpatterns that must be explored along the search for a maximal frequent pattern. In this paper a fast probabilistic hill climbing search that optimizes a pattern interest function is used to rapidly discover interesting feature set patterns.

Phrase 1														
int	$\boxed{+2}$	$\boxed{+1}$	$\boxed{+2}$	$+2$	$+3$	-5	-3	$+1$	$\boxed{+2}$	$+3$	$\boxed{-1}$	$\boxed{-2}$	-2	-1
c5	$+s$	$+s$	$+s$	$+s$	$+l$	$-l$	$-l$	$+s$	$+s$	$+l$	$-s$	$-s$	$-s$	$-s$
c3	u	u	u	u	u	d	d	u	u	u	d	d	d	d
Phrase 2														
int	$\boxed{+2}$	$\boxed{+1}$	$\boxed{+2}$	-2	-1	-2	-2	$+2$	$\boxed{+2}$	$+1$	$\boxed{-1}$	$\boxed{-2}$	-1	$+1$
c5	$+s$	$+s$	$+s$	$-s$	$-s$	$-s$	$-s$	$+s$	$+s$	$+s$	$-s$	$-s$	$-s$	$+s$
c3	u	u	u	d	d	d	d	u	u	u	d	d	d	u
Pattern	$\{s,u\}$	$\{s,u\}$	$\{s,u\}$	$\{s\}$	$\{\}$	$\{d\}$	$\{d\}$	$\{s,u\}$	$\{s,u\}$	$\{u\}$	$\{s,d\}$	$\{s,d\}$	$\{s,d\}$	$\{s\}$

Fig. 1. Top: first two phrases of Brassens' *Le Bistrot*, encoded using three viewpoints of different abstraction levels (middle). Bottom: the most specific feature set pattern that occurs in both phrases using the taxonomy of Fig. 2. Conserved melodic intervals are boxed.

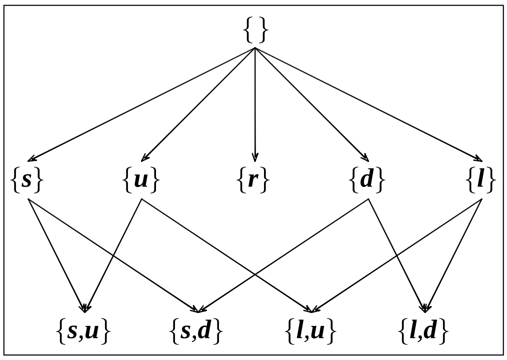

Fig. 2. Taxonomy of satisfiable generalized melodic intervals represented as feature sets.

A. Viewpoints, features, and feature sets

A *viewpoint* is a partial function which computes a *viewpoint element* for events in a sequence [6]. For example, for the melodic interval viewpoint ***int***, the viewpoint elements represent the difference in semitones between a note and its preceding note. A viewpoint/element pair is called a *feature*. A set of features is called a *feature set*.

The generalized interval taxonomy in Figure 2 uses the five viewpoints ***u*** (up), ***d*** (down), ***r*** (repeat), ***s*** (step), and ***l*** (leap). As the elements of these viewpoints are either undefined or true for notes, it is convenient to write only the name of the viewpoint and call such viewpoint names *classes*. Negated interval classes (e.g., "not a step") can also easily be handled in this scheme, but are not the focus of this paper.

A sequence of feature sets is called a *feature set pattern*. In a corpus, the *total count* of a pattern is the number of positions where it occurs (not counting overlapping occurrences). The *piece count* of a pattern is the number of pieces that have one or more occurrences of the pattern. A *frequent pattern* is one that occurs with at least a specified minimum total count and piece count.

B. Pattern interest

Patterns can be ranked according to the difference between observed and expected counts in a corpus ([5], [7]). Large differences indicate potentially interesting patterns. Here, a pattern P is given a *pattern interest* $I(P)$, which represents the ratio of observed to expected total counts:

$$I(P) = \frac{C_t(P)}{E_t(P)}, \tag{1}$$

where $C_t(\cdot)$ represents the total count in a corpus and $E_t(\cdot)$ expected total count.

To define the expected total count, consider a feature set pattern $P = f_1, \ldots, f_n$. The (non-overlapping) total count of the pattern can be no more than $C_t(\{\})/n$. The expected total count in the corpus is this quantity multiplied by the probability of the pattern:

$$E_t(P) = \frac{C_t(\{\})}{n} \times p(f_1, \ldots, f_n).$$

The probability of the pattern is simply the product of the relative frequencies of component feature sets within the corpus:

$$p(f_1, \ldots, f_n) = \prod_{i=1}^{n} \frac{C_t(f_i)}{C_t(\{\})}.$$

The pattern interest $I(\cdot)$ can be computed rapidly since the feature set counts emerge directly from the first phase of the discovery algorithm, as described below.

C. Searching for interesting patterns

A corpus (which may be one, two, or any number of pieces) is first transformed to complete feature set sequences by *saturation*: applying each viewpoint in a catalog to every note in every piece. The pattern discovery algorithm then proceeds in two phases. In the first phase, all frequent feature sets are found, and configured into a subsumption taxonomy using a description logic classification algorithm [2] which places each feature set

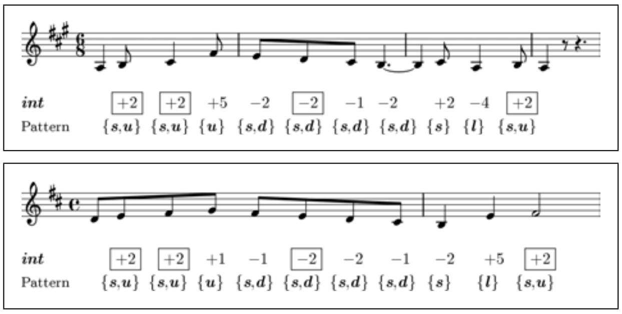

Fig. 4. Excerpts from Brassens' *Les quatre bacheliers* (phrase 2, top), and *Rien à jeter* (phrase 3, bottom), showing a discovered generalized interval pattern.

just below all most specific subsumers and just above all most general subsumees.

The second phase of the discovery algorithm uses this taxonomy in the exploration of the space of frequent feature set patterns. This space can be fully visited by recursively applying two specialization operators [1] beginning at the empty pattern: an I-step which specializes the right-most component of a pattern by walking one step down the feature set subsumption taxonomy (there may be several such specializations); and an S-step which appends the empty feature set $\{\}$ to a pattern.

The search space of frequent patterns can be huge, even for a piece or corpus of moderate size, and it is necessary to prune the space to arrive at a solution quickly. Here a heuristic probabilistic hill climbing method is developed. The interest measure (Expression 1) is calculated for all candidate specializations (all possible I-steps, and the S-step) and one of these is sampled. This method is similar to the Gibbs sampling methods used for protein motif discovery [9]. The probability of sampling a pattern P from the set Q of all candidate specializations is

$$\frac{I(P)}{\sum_{q \in Q} I(q)},$$

and the sampled pattern then becomes the new current pattern. The process is started at the empty pattern, and iterated until there are no further candidate specializations (i.e., none of them are frequent).

To reduce the effect of climbing to solutions with low pattern interest, a random restart method is employed. The probabilistic hill climbing search is restarted 10 times, and the best pattern arising from all iterations is reported.

III. Results

A corpus of 115 Brassens songs (average 80 notes per song) in the MIDI format was transformed to feature set sequences using the six viewpoints ***u***, ***d***, ***r***, ***l***, ***s***, and ***int***, and a minimum total count of 2. Two types of analyses were performed using the pattern discovery method described above: intraopus analysis (patterns repeated in a single piece), and comparative analysis (patterns found in two pieces). In both experiments, patterns were restricted to include at most one empty feature set.

In this section a few results of the method are presented. The patterns in Figures 1 and 5 were both discovered by intraopus analysis. For comparative analysis, all pairs of pieces from the corpus were provided as input to the pattern discovery method, and several interesting patterns emerged.

Phrase 1					
int	-1	$+1$	$+4$	-12	$+4$
c5	$-s$	$+s$	$+l$	$-l$	$+l$
Phrase 2					
int	-5	$+5$	$+3$	-3	-5
c5	$-l$	$+l$	$+l$	$-l$	$-l$
Pattern	$\{d\}$	$\{u\}$	$\{l,u\}$	$\{l,d\}$	$\{l\}$

Fig. 5. Top: endings of the first two phrases of Brassens' *Tonton Nestor*, encoded using the ***int*** and ***c5*** viewpoints (middle). Bottom: the most specific pattern that captures both phrases according to the generalized melodic interval taxonomy of Figure 2.

Figure 3, for example, illustrates a fragment of a pattern discovered in Brassens' *Embrasse-les tous* and *Stances à un cambrioleur*. Intervals shared by the melodies are not explicitly shown in the pattern, but rather shown in boxes. The pattern has clearly captured the similarity in arpeggiation over a triad; major then minor in one piece and diminished in the other. The pattern is unique to these two pieces, not found anywhere else in the corpus. The results also illustrate an unwelcome aspect of the probabilistic hill climbing method: after correctly spanning the first 6 arpeggiations, the pattern continues to be inappropriately extended with a few general feature sets because they continue to increase the interest measure of the pattern. This is a topic of current research.

In Figure 4, two phrases from Brassens' *Les quatre bacheliers* and *Rien à jeter* are shown, along with a fragment of a discovered feature set pattern (again, conserved intervals are boxed rather than written in the pattern). Both pieces have a similar overall form: four phrases, each phrase having 8 initial notes in a convex shape, followed by 3 to 5 notes and a cadence. The pattern spans two full phrases (phrases 2 and 3 of *Les...*, and phrases 3 and 4 of *Rien...*) and reveals their similarity in melodic shape. This convex shape (first 7 components of the pattern) occurs in a total of 13 Brassens songs.

IV. Discussion

This paper has described the use of generalized interval classes combined with a fast pattern discovery algorithm. The method was applied to the music of French singer-songwriter Georges Brassens, discovering patterns in single pieces and pairs of pieces.

The use of feature sets to represent melodic contour classes was inspired by the work of Cambouropoulos et al. [4], who add two ***c3*** classes u (up), and d (down) to the melodic contour ***c5***. Their method discovers repe-

Fig. 3. A fragment of a pattern discovered in Brassens' *Embrasse-les tous* (top) and *Stances à un Cambrioleur* (bottom). The pattern extends further to the left and the whole pattern captures the first three phrases of both songs.

titions in a single melody by first transforming it into a string of extended ***c5*** identifiers, then applying a modified suffix tree method. This preprocessing will transform the melodic intervals of $+2$ and -2 into the classes ***u*** and ***d***, and special matching rules ensure that, for example, ***u*** matches either a step up ($+s$) or a leap up ($+l$). This preprocessing leads to incompleteness in the pattern discovery phase: for example, a melodic interval of $+1$ will never be aligned with an interval of $+3$, even though they both belong to the class ***u***. Figure 5 illustrates this problem with a music example. The contextual similarity of intervals -1 and -5 (note 2) and $+1$ and $+5$ (note 3) cannot be discovered. Furthermore, to represent the contextual similarity of intervals $+4$ and -5 (note 6) the general interval class ***l*** (leap), which is not accommodated by their method, is required.

The algorithm described in this paper directly uses a subsumption taxonomy of generalized intervals represented by feature sets. This means that the ***s*** (step) and ***l*** (leap) interval classes (without regard to direction of motion) can be easily captured and that no single level of abstraction needs to be specified in advance. The expressive representation for pattern components reported here, combined with a fast probabilistic hill climbing algorithm for pattern discovery, can produce interesting musical results.

REFERENCES

[1] M. Bergeron and D. Conklin, "Representation and discovery of feature set patterns in music," in *International Workshop on Artificial Intelligence and Music, at IJCAI-07, The Twentieth International Joint Conference on Artificial Intelligence*, 2007, pp. 1–12.

[2] R. Brachman and H. Levesque, *Knowledge Representation and Reasoning*. Morgan Kaufmann, 2004.

[3] G. Brassens, *Poèmes & chansons*. Éditions du Seuil, 1993.

[4] E. Cambouropoulos, M. Crochemore, C. Iliopoulos, M. Mohamed, and M.-F. Sagot, "A pattern extraction algorithm for abstract melodic representations that allow partial overlapping of intervallic categories," in *International Conference on Music Information Retrieval (ISMIR)*, 2005, pp. 167–174.

[5] D. Conklin and C. Anagnostopoulou, "Representation and discovery of multiple viewpoint patterns," in *International Computer Music Conference (ICMC)*, 2001, pp. 479–485.

[6] D. Conklin, "Melodic analysis with segment classes," *Machine Learning*, vol. 65, no. 2–3, pp. 349–360, 2006.

[7] D. Huron, "What is a musical feature? Forte's analysis of Brahms's Opus 51, No. 1, revisited," *Music Theory Online*, vol. 7, no. 4, 2001.

[8] O. Lartillot, "A musical pattern discovery system founded on a modeling of listening strategies," *Computer Music Journal*, vol. 28, no. 3, pp. 53–67, 2004.

[9] C. Lawrence, S. Altschul, M. Boguski, J. Liu, A. Neuwald, and J. Wootton, "Detecting subtle sequence signals: a Gibbs sampling strategy for multiple alignment," *Science*, vol. 262, no. 5131, pp. 208–214, 1993.